

HYPERBOLIC AND PARABOLIC UNIMODULAR RANDOM MAPS

OMER ANGEL TOM HUTCHCROFT ASAF NACHMIAS GOURAB RAY

ABSTRACT. We show that for infinite planar unimodular random rooted maps, many global geometric and probabilistic properties are equivalent, and are determined by a natural, local notion of average curvature. This dichotomy includes properties relating to amenability, conformal geometry, random walks, uniform and minimal spanning forests, and Bernoulli bond percolation. We also prove that every simply connected unimodular random rooted map is sofic, that is, a Benjamini-Schramm limit of finite maps.

1 Introduction

In the classical theory of Riemann surfaces, the Uniformization Theorem states that every simply connected, non-compact Riemann surface is conformally equivalent to either the plane or the disc, which are inequivalent to each other by Liouville's theorem. The dichotomy provided by this theorem manifests itself in several different ways, relating to analytic, geometric and probabilistic properties of surfaces. In particular, if S is a simply connected, non-compact Riemann surface, then either

S is parabolic: it is conformally equivalent to the plane, admits a compatible Riemannian metric of constant curvature 0, does not admit non-constant bounded harmonic functions, and is recurrent for Brownian motion,

or else

S is hyperbolic: it is conformally equivalent to the disc, admits a compatible Riemannian metric of constant curvature -1 , admits non-constant bounded harmonic functions, and is transient for Brownian motion.

In the 1990's, a discrete counterpart to this dichotomy began to develop in the setting of bounded degree planar graphs [8, 44, 43, 18, 19]. A major milestone in this theory was the seminal work of He and Schramm [44, 43], who studied *circle packings* of infinite triangulations of the plane. They proved that every infinite triangulation of the plane can be circle packed in either the unit disc or in the plane, but not both. A triangulation is called **CP hyperbolic** or **CP parabolic** accordingly. He and Schramm also connected the circle packing type to isoperimetric and probabilistic properties of the triangulation, showing in particular that, in the bounded degree case, CP parabolicity is equivalent to the recurrence of simple random walk. Later, Benjamini and Schramm [18, 19] provided an analytic aspect to this dichotomy, showing that every bounded degree, infinite planar graph admits non-constant bounded harmonic functions if and only if it is transient for simple random walk, and that in this case the graph also admits non-constant bounded harmonic functions of finite Dirichlet energy. Most of this theory fails without the assumption of bounded degrees, as one can easily construct pathological counter-examples to the theorems above.

The goal of this paper is to develop a similar theory for *unimodular random rooted maps*, without the assumption of bounded degree. In our earlier work [4], we studied circle packings

of, and random walks on, random plane triangulations of unbounded degree. In this paper, we study many further properties of unimodular random planar maps, which we do not assume to be triangulations. Our results also have consequences for unimodular random maps that are not planar, which we develop in [Section 7](#). Our main result may be stated informally as follows; see [Theorem 1](#) for a complete and precise statement.

Theorem (The Dichotomy Theorem). *Every infinite, planar, unimodular random rooted planar map is either hyperbolic or parabolic. The map is hyperbolic if and only if its average curvature is negative and is parabolic if and only if its average curvature is zero. The type of a unimodular random rooted map determines many of its properties.*

The many properties we show to be determined by the type of the map are far-reaching, relating to aspects of the map including amenability, random walks, harmonic functions, spanning forests, Bernoulli bond percolation, and the conformal type of associated Riemann surfaces. The seeds of such a dichotomy were already apparent in [4], in which we proved that a unimodular random rooted plane triangulation is CP parabolic almost surely if and only if the expected degree of the root is six (which is equivalent to the average curvature being zero), if and only if the triangulation is **invariantly amenable** – a notion of amenability due to Aldous and Lyons [2] that is particularly suitable to unimodular random rooted graphs. A notable property that is *not* a part of [Theorem 1](#) is recurrence of the random walk: while every hyperbolic unimodular random rooted map is transient, not every parabolic unimodular random rooted map is recurrent. ([Theorem 1](#) can be combined with the work of Gurel-Gurevich and the third author [38] to deduce that a parabolic unimodular random planar map is recurrent under the additional assumption that the degree of the root has an *exponential tail*.)

A **map** is a *proper* (i.e., locally finite) embedding of a graph into an oriented surface viewed up to orientation preserving homeomorphisms of the surface. (Other definitions extend to non-orientable maps, but we shall not be concerned with those here.) A **rooted map** is a map together with a distinguished root vertex. The map is called **planar** if the surface is homeomorphic to an open subset of the sphere, and is **simply connected** if the surface is homeomorphic to the sphere or the plane. (In particular, every simply connected map is planar.) A random rooted map is said to be **unimodular** if it satisfies the **mass-transport principle**, which can be interpreted as meaning that ‘every vertex of the map is equally likely to be the root’. See [Section 2](#) for precise definitions of each of these terms.

The **curvature** of a map is a local geometric property, closely related to the Gaussian curvature of manifolds that may be constructed from the map; see [Section 4](#) for a precise definition. For one natural manifold constructed from the map by gluing together regular polygons, the curvature at each vertex v is

$$\kappa(v) = 2\pi - \sum_{f \perp v} \frac{\deg(f) - 2}{\deg(f)} \pi,$$

where the sum is taken over faces of the map incident to v , and a face is counted with multiplicity if more than one of the corners of the face are located at v . We can define the average curvature of (M, ρ) to be the expectation $\mathbb{E}[\kappa(\rho)]$. [Theorem 4.10](#) states that the average curvature is a canonical quantity associated to the random map, in the sense that any unimodular way of associating a manifold to the map will result in the same average curvature. Observe that the average curvature of a unimodular random triangulation is equal

to $(6 - \mathbb{E}[\deg(\rho)])\pi/3$, so that a unimodular triangulation has expected degree greater than six if and only if it has negative average curvature. This relates the dichotomy described in [4] to that of [Theorem 1](#).

Classical examples of unimodular random maps are provided by Voronoi diagrams of stationary point processes [17] and (slightly modified) Galton-Watson trees [2, Example 1.1], as well as lattices in the Euclidean and hyperbolic planes, and arbitrary local limits of finite maps. Many local modifications of maps, such as taking Bernoulli percolation or uniform or minimal spanning trees, preserve unimodularity, giving rise to many additional examples. Unimodular random maps, most notably the uniform infinite planar triangulation (UIPT) [7] and quadrangulation (UIPQ) [54, 29], have also been studied in the context of 2-dimensional quantum gravity; see the survey [35] and references therein. More recently, hyperbolic variants of the UIPQ have been constructed [5, 28]. Many of these examples do not have uniformly bounded degrees, so that the deterministic theory is not applicable to them.

1.1 The Dichotomy Theorem

Since many of the notions tied together in the following theorem are well known, we first state the theorem, and defer definitions to individual sections dealing with each of the properties.

Theorem 1 (The Dichotomy Theorem). *Let (M, ρ) be an infinite, ergodic, unimodular random rooted planar map and suppose that $\mathbb{E}[\deg(\rho)] < \infty$. Then the average curvature of (M, ρ) is non-positive and the following are equivalent:*

- (1) (M, ρ) has average curvature zero.
- (2) (M, ρ) is invariantly amenable.
- (3) Every bounded degree subgraph of M is amenable almost surely.
- (4) Every subtree of M is amenable almost surely.
- (5) Every bounded degree subgraph of M is recurrent almost surely.
- (6) Every subtree of M is recurrent almost surely.
- (7) (M, ρ) is a Benjamini-Schramm limit of finite planar maps.
- (8) (M, ρ) is a Benjamini-Schramm limit of a sequence $\langle M_n \rangle_{n \geq 0}$ of finite maps such that

$$\frac{\text{genus}(M_n)}{\#\{\text{vertices of } M_n\}} \xrightarrow{n \rightarrow \infty} 0.$$

- (9) The Riemann surface associated to M is conformally equivalent to either the plane \mathbb{C} or the cylinder \mathbb{C}/\mathbb{Z} almost surely.
- (10) M does not admit any non-constant bounded harmonic functions almost surely.
- (11) M does not admit any non-constant harmonic functions of finite Dirichlet energy almost surely.
- (12) The laws of the free and wired uniform spanning forests of M coincide almost surely.
- (13) The wired uniform spanning forest of M is connected almost surely.
- (14) Two independent random walks on M intersect infinitely often almost surely.
- (15) The laws of the free and wired minimal spanning forests of M coincide almost surely.
- (16) Bernoulli(p) bond percolation on M has at most one infinite connected component for every $p \in [0, 1]$ almost surely (in particular, $p_c = p_u$).
- (17) M is vertex extremal length parabolic almost surely.

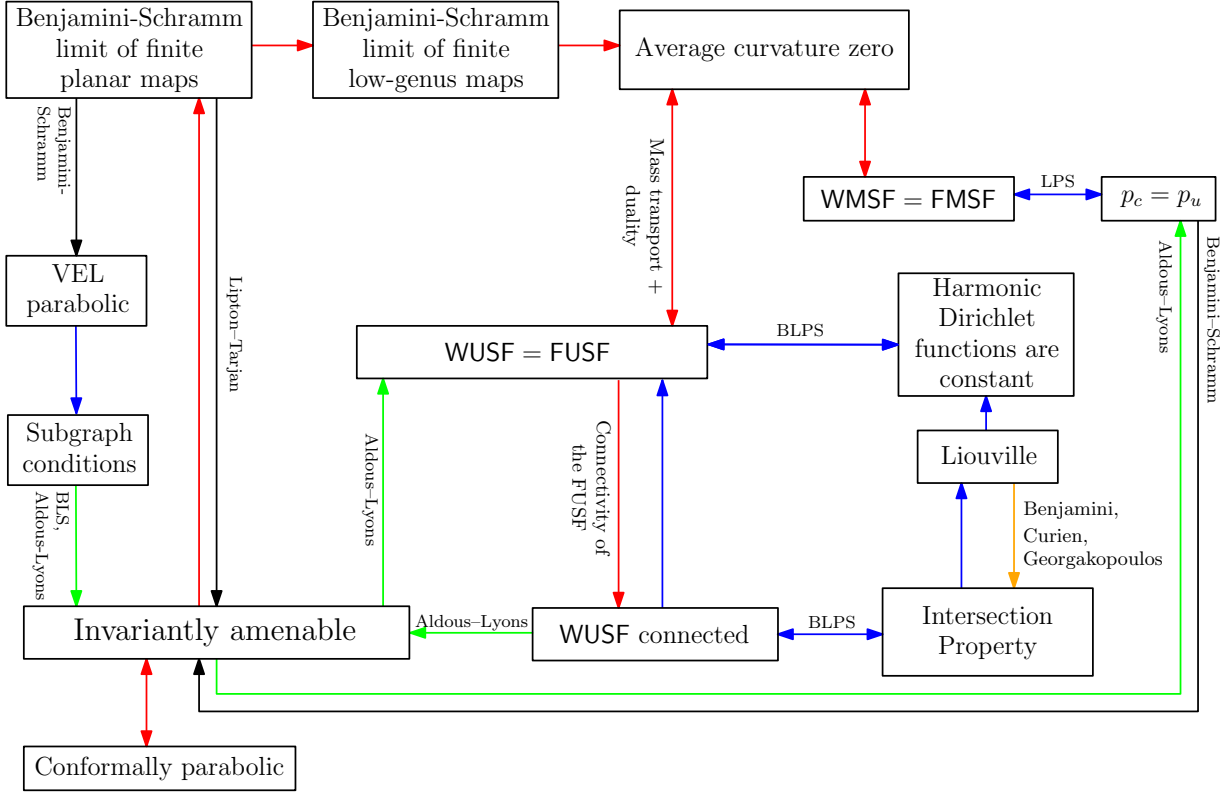



Figure 1. The logical structure for the proof of [Theorem 1](#) in the simply connected case. Implications new to this paper are in red. Blue implications hold for arbitrary graphs; the orange implication holds for arbitrary planar graphs, and green implications hold for unimodular random rooted graphs even without planarity. A few implications between items that are known but not used in the proof are omitted.

In light of this theorem, we call a unimodular random rooted map (M, ρ) with $\mathbb{E}[\deg(\rho)] < \infty$ **parabolic** if its average curvature is zero (and, in the planar case, clauses (1)–(17) all hold), and **hyperbolic** if its average curvature is negative (and, in the planar case, the clauses all fail).

As one might guess from the structure of [Theorem 1](#), the proof consists of many separate arguments for the different implications. Some of the implications are already present in the literature, and part of this paper is spent surveying the earlier works that form the individual implications between the long list of equivalent items in [Theorem 1](#). For the sake of completeness we also include proofs of several standard technical results from the ergodic theory literature using probabilistic terminology.

Some of the implications in [Theorem 1](#) hold in any graph. For example, (10) implies (11) for any graph, and (14) and (13) are always equivalent (see [14]). Other implications hold for any planar graph. For example (10) is equivalent to (14), see [11]. We do not provide a comprehensive list of the assumptions needed for each implication, but some of this information is encoded in [Figure 1](#).



Most of the paper is dedicated to proving the theorem under the additional assumption that M is simply connected; the multiply-connected case is easier and is handled separately in [Section 7](#). The logical structure of the proof in the simply connected case is summarized in [Figures 1](#) and [2](#). [Figure 1](#) also shows which implications were already known and which are proved in the present paper.

1.2 Unimodular planar maps are sofic

It is easy to see that every (possibly random) finite graph with conditionally uniform root is unimodular. Moreover, unimodularity is preserved under distributional limits in the local topology. It follows that every Benjamini-Schramm limit of finite random graphs is unimodular. A random rooted graph that can be obtained in this way is called **sofic**. It is a major open problem to determine whether the converse holds, that is, whether every

unimodular random rooted graph is sofic [2, Section 10]. The next theorem answers this question positively for simply connected unimodular random maps.

Theorem 2. *Every simply connected unimodular random rooted map is sofic.*

The proof of [Theorem 2](#) relies on the corresponding result for trees, which is due to Bowen [24], Elek [32], and Benjamini, Lyons and Schramm [16]. As was observed by Elek and Lippner [33], it follows that *treeable* unimodular graphs (that is, unimodular graphs that exhibit an invariant spanning tree) are sofic, and so the key new step is proving the connectivity of the free uniform spanning forest.

Note that the finite maps converging to a given infinite unimodular random rooted map (M, ρ) need not be planar. Indeed, [Theorem 1](#) characterises the Benjamini-Schramm limits of finite *planar* maps exactly as the parabolic unimodular random rooted maps. Moreover, if $\langle M_n \rangle_{n \geq 0}$ is a sequence of finite maps converging to an infinite hyperbolic unimodular random rooted map, then the approximating maps M_n must have genus comparable to their number of vertices as n tends to infinity.

2 Unimodular Maps

2.1 Maps

We provide here a brief background to the concept of maps, and refer the reader to [55, Chapter 1.3] for a comprehensive treatment. Let $G = (V, E)$ be a connected graph, which may contain self-loops and multiple edges. An **embedding** of a graph in a surface S is a drawing of the graph in the surface with non-crossing edges. Given an embedding, the connected components of the complement of the image of G are called **faces**. An embedding is said to be **proper** if the following conditions hold:

- (1) it is locally finite (every compact set in S intersects finitely many edges),
- (2) every face is homeomorphic to an open disc, and
- (3) for every face f , if we consider the oriented edges of G that have their right hand side incident to f , and consider the permutation that maps each such oriented edge to the oriented edge following it in the clockwise order around the face, then this permutation has a single orbit.

For example, the complete graph on three vertices can be properly embedded in the sphere but not in the plane. If S is simply connected or compact, then any embedding that satisfies (1) and (2) must also satisfy (3). For this reason, the condition (3) is not included in many references that deal primarily with finite maps. An example of an embedding that satisfies (1) and (2) but not (3) is given by drawing \mathbb{Z} along a straight line in an infinite cylinder.

We define a (locally finite) **map** M to be a connected, locally finite graph G together with an equivalence class of proper embeddings of G into oriented surfaces, where two embeddings are equivalent if there is an orientation preserving homeomorphism between the two surfaces that sends one embedding to the other. If M is a map with underlying graph G , we refer to any proper embedding of G that falls into the equivalence class of embeddings corresponding to M as an embedding of M . A map is said to be **planar** if it is embedded into a surface homeomorphic to an open subset of the sphere, and is said to be **simply connected** if it is embedded into a simply connected surface (which is necessarily homeomorphic to either the sphere or the plane).

Let M be a map with underlying graph G and let z be a proper embedding of M into a surface S . If every face of M has finite degree, the **dual map** of M , denoted M^\dagger is defined as follows. The underlying graph of M^\dagger , denoted G^\dagger , has the faces of M as vertices, and has an edge drawn between two faces of M for each edge in M that is incident to both of the faces. We define an embedding z^\dagger of G^\dagger into S by placing each vertex of G^\dagger in the interior of the corresponding face of M and each edge of G^\dagger so that it crosses the corresponding edge of G but no others. We define M^\dagger to be the map with underlying graph G represented by the pair (S, z^\dagger) : Although the embedding z^\dagger is not uniquely defined, every choice of z^\dagger defines the same map. The construction gives a canonical bijection between edges of G and edges of G^\dagger . We write e^\dagger for the edge of G^\dagger corresponding to e . If e is an oriented edge, we let e^\dagger be oriented so that it crosses e from right to left as viewed from the orientation of e .

Despite their topological definitions, maps and their duals can in fact be defined entirely combinatorially. Given any graph, we consider each edge as two oriented edges in opposite directions. We write $E^\rightarrow = E^\rightarrow(G)$ for the set of oriented edges of a graph. For each directed edge e , we have a head e^+ and tail e^- , and write $-e$ for the reversal of e . Given a map M and a vertex v of M , let $\sigma_v = \sigma_v(M)$ be the cyclic permutation of the set $\{e \in E^\rightarrow : e^- = v\}$ of oriented edges emanating from v corresponding to counter-clockwise rotation in S . This procedure defines a bijection between maps and graphs labelled by cyclic permutations.

Theorem 2.1 ([55]). *Given a connected, locally finite graph G and a collection of cyclic permutations σ_v of the sets $\{e \in E^\rightarrow : e^- = v\}$, there exists a unique map M with underlying graph G such that $\sigma(M) = \sigma$.*

In light of [Theorem 2.1](#), we formally identify a map M with the pair (G, σ) . Given such a combinatorial specification of a map M as a pair (G, σ) , we may form an embedding of the map into a surface $S(M)$ by gluing topological discs according to the combinatorics of the map (see [Figure 3](#)).

Note that we can write $\sigma(e)$ for $\sigma_v(e)$, since necessarily $v = e^-$. Thus σ is a permutation on the set of directed edges of M . Formally, a **corner** in the map at a vertex v is an ordered pair of directed edges $(e, \sigma(e))$, with $e^- = v$. Of course, a corner is determined by the directed edge e . Just as vertices of M are orbits of σ , the faces of a map $M = (G, \sigma)$ can be defined abstractly as orbits of the permutation $\sigma^\dagger : E^\rightarrow \rightarrow E^\rightarrow$ defined by $\sigma^\dagger(e) = \sigma^{-1}(-e)$ for each $e \in E^\rightarrow$. The dual e^\dagger of a directed edge e is defined to have the orbit of e as its tail and the orbit of $-e$ as its head, so that we again have a bijection between directed edges of M and their duals. Using this bijection, we can consider σ^\dagger to acts on dual edges, and the dual map M^\dagger is then constructed abstractly as $M^\dagger = (G^\dagger, \sigma^\dagger)$. We define maps with infinite degree vertices, and duals of maps with infinite degree faces, directly through this abstract formalism.

If e is an oriented edge in a map, we write e^ℓ for the face of M to the left of e and e^r for the face of M to the right of e , so that $e^\ell = (e^\dagger)^-$ and $e^r = (e^\dagger)^+$. Given a map M , we write $f \perp v$ if the face f is incident to the vertex v , that is, if there exists an oriented edge e of M such that $e^- = v$ and $e^\ell = f$. When writing a sum of the form $\sum_{f \perp v}$, we use the convention that a face f is counted with multiplicity according to the number of oriented edges e of M such that $e^- = v$ and $e^\ell = f$. Similarly, when writing a sum of the form $\sum_{u \sim v}$, we count each vertex u with multiplicity according to the number of oriented edges e of M such that $e^- = v$ and $e^+ = u$. The degree of a face is defined to be the number of oriented edges with $e^r = f$, i.e., the degree of the face in the dual.




Figure 3. Different maps with the same underlying graph. The two maps both have K_4 as their underlying graph, but the left is a sphere map while the right is a torus map. Both maps are represented both abstractly as a graph together with a cyclic permutation of the edges emanating from each vertex (left) and as a graph embedded in a surface (right).

2.2 Unimodularity and the mass transport principle

As noted, a **rooted graph** (G, ρ) is a connected, locally finite (multi)graph $G = (V, E)$ together with a distinguished vertex ρ , called the **root**. A graph isomorphism $\phi : G \rightarrow G'$ is an isomorphism of rooted graphs is a graph isomorphism that maps the root to the root.

The **local topology** (see [21]) is the topology on the set \mathcal{G}_\bullet of isomorphism classes of rooted graphs induced by the metric

$$d_{\text{loc}}((G, \rho), (G', \rho')) = e^{-R},$$

where

$$R = R((G, \rho), (G', \rho')) = \sup \{R \geq 0 : B_R(G, \rho) \cong B_R(G', \rho')\},$$

i.e., the maximal radius such that the balls $B_R(G, \rho)$ and $B_R(G', \rho')$ are isomorphic as rooted graphs.

A **random rooted graph** is a random variable taking values in the space \mathcal{G}_\bullet endowed with the local topology. Similarly, a **doubly-rooted graph** is a graph together with an ordered pair of distinguished (not necessarily distinct) vertices. Denote the space of isomorphism classes of doubly-rooted graphs equipped with this topology by $\mathcal{G}_{\bullet\bullet}$.

The spaces of rooted and doubly-rooted maps are defined similarly and are denoted \mathcal{M}_\bullet and $\mathcal{M}_{\bullet\bullet}$ respectively. For this, an isomorphism ϕ of rooted maps is preserves the roots and the map structure, so that $\phi \circ \sigma = \sigma' \circ \phi$ for vertices inside the balls.

A further generalisation of these spaces will also be useful. A **marked graph** (referred to by Aldous and Lyons [2] as a network) is defined to be a locally finite, connected graph together with a function $m : E \cup V \rightarrow \mathbb{X}$ assigning each vertex and edge of G a **mark** in some Polish space \mathbb{X} , referred to as the **mark space**. The local topology on the set of isomorphism classes of rooted graphs with marks in \mathbb{X} is the topology induced by the metric

$$d_{\text{loc}}((G, \rho, m), (G', \rho', m')) = e^{-R}$$

where $R((G, \rho, m), (G', \rho', m'))$ is the largest R such that there exists an isomorphism of rooted graphs $\phi : (B_G(\rho, b), m, \rho) \rightarrow (B_{G'}(\rho', b), m', \rho')$ such that $d_{\mathbb{X}}(m'(\phi(x)), m(x)) \leq 1/R$ for every vertex or edge x of $B_G(\rho, n)$, and $d_{\mathbb{X}}$ is a metric compatible with the topology of \mathbb{X} . The space of isomorphism classes of rooted marked graphs with marks in \mathbb{X} is denoted $\mathcal{G}_\bullet^{\mathbb{X}}$. The space of rooted marked maps is defined similarly.

It is possible to consider rooted maps and rooted marked maps as rooted marked graphs by encoding the permutations σ_v in the marks – in particular, this means that any statement

that holds for all unimodular random rooted marked graphs also holds for all unimodular random rooted marked maps. See [2, Example 9.6].

A **mass transport** is a Borel function $f : \mathcal{G}_{\bullet\bullet} \rightarrow [0, \infty]$. A random rooted graph (G, ρ) is said to be **unimodular** if it satisfies the **Mass Transport Principle**: for every mass transport f ,

$$\mathbb{E} \left[\sum_{v \in V} f(G, \rho, v) \right] = \mathbb{E} \left[\sum_{u \in V} f(G, u, \rho) \right]. \quad (\text{MTP})$$

That is,

$$\textit{Expected mass out equals expected mass in.}$$

Unimodularity of random rooted maps, marked graphs and marked maps are defined similarly. A probability measure \mathbb{P} on \mathcal{G}_{\bullet} is said to be unimodular if a random rooted map with law \mathbb{P} is unimodular. Unimodular probability measures on rooted maps, marked graphs and marked maps are defined similarly.

It will also be convenient to work with signed mass transports. If a measurable $f : \mathcal{G}_{\bullet\bullet} \rightarrow \mathbb{R}$ satisfies

$$\mathbb{E} \left[\sum_{v \in V} |f(G, \rho, v)| \right] < \infty,$$

then the conclusion of the mass transport principle holds for f . This follows easily from the usual mass transport principle for non-negative f .

The mass transport principle was first introduced by Häggström [40] to study dependent percolation on Cayley graphs. The formulation of the mass transport principle presented here was suggested by Benjamini and Schramm [21] and developed systematically by Aldous and Lyons [2]. The unimodular probability measures on \mathcal{G}_{\bullet} form a weakly closed, convex subset of the space of probability measures on \mathcal{G}_{\bullet} , so that weak limits of unimodular random graphs are unimodular. In particular, a weak limit of finite graphs with uniformly chosen roots is unimodular: such a limit of finite graphs is referred to as a **Benjamini-Schramm limit**. It is a major open problem to determine whether all unimodular random rooted graphs arise as Benjamini-Schramm limits of finite graphs [2, §10].

We will make frequent use of the fact that unimodularity is stable under most reasonable ways of modifying a graph locally without changing its vertex set. The following lemma is a formalization of this fact.

Lemma 2.2. *Let \mathbb{X}_1 and \mathbb{X}_2 be polish spaces. Let (G, ρ, m) be a unimodular random rooted \mathbb{X}_1 -marked graph, and suppose that (G', ρ, m') is a random rooted marked graph with the same vertex set as G , such that, for every pair of vertices u, v in G , the conditional distribution of (G', u, v, m') given (G, ρ, m) coincides a.s. with some measurable function of the isomorphism class of (G, u, v, m) . Then (G', ρ, m') is unimodular.*

Proof. Let $f : \mathcal{G}_{\bullet\bullet}^{\mathbb{X}_2} \rightarrow [0, \infty]$ be a mass transport. For each pair of vertices u, v in G , the conditional expectation

$$F(G, u, v, m) = \mathbb{E} [f(G', u, v, m') \mid (G, \rho, m)]$$

coincides a.s. with a measurable function of (G, u, v, m) , and so is itself a mass transport. Since (G, ρ, m) is unimodular we deduce that

$$\begin{aligned} \mathbb{E} \sum_v f(G', \rho, v, m') &= \mathbb{E} \sum_v F(G, \rho, v, m) \\ &= \mathbb{E} \sum_v F(G, v, \rho, m) = \mathbb{E} \sum_u f(G', u, \rho, m'), \end{aligned} \quad (2.1)$$

verifying that (G', ρ, m') satisfies the mass transport principle. \square

Later in the paper, we will often claim that various random rooted graphs obtained from unimodular random rooted graphs in various ways are unimodular. These claims can always be justified either as a direct consequence of [Lemma 2.2](#), or otherwise by applying the mass transport principle to the conditional expectation as in [eq. \(2.1\)](#).

2.3 Reversibility

Recall that the **simple random walk** on a locally finite graph $G = (V, E)$ is the Markov process $\langle X_n \rangle_{n \geq 0}$ on the state space V with transition probabilities $p(u, v)$ defined to be the fraction of edges emanating from u that end in v . A probability measure \mathbb{P} on \mathcal{G}_\bullet is said to be **stationary** if, when (G, ρ) is a random rooted graph with law \mathbb{P} and $\langle X_n \rangle_{n \geq 0}$ is a simple random walk on G started at the root, (G, ρ) and (G, X_n) have the same distribution for all n . The measure \mathbb{P} is said to be **reversible** if furthermore

$$(G, \rho, X_n) \stackrel{d}{=} (G, X_n, \rho)$$

for all n . A random rooted graph is said to be reversible if its law is reversible. Reversible random rooted maps, marked graphs and marked maps are defined similarly. Every reversible random rooted graph is clearly stationary, but the converse need not hold in general [13, Examples 3.1 and 3.2].

Let $\mathcal{G}_{\leftrightarrow}$ be the space of isomorphism classes of connected, locally finite graphs equipped with a distinguished bi-infinite path, which we endow with a natural variant of the local topology. If (G, ρ) is a reversible random rooted graph and $\langle X_n \rangle_{n \geq 0}$ and $\langle X_{-n} \rangle_{n \geq 0}$ are independent simple random walks started from ρ , then the sequence

$$\langle (G, \langle X_{n+k} \rangle_{n \in \mathbb{Z}}) \rangle_{k \in \mathbb{Z}}$$

of $\mathcal{G}_{\leftrightarrow}$ -valued random variables is stationary.

The following correspondence between unimodular and reversible random rooted graphs is implicit in [2] and was proven explicitly in [9]. Similar correspondences also hold between unimodular and reversible random rooted maps, marked graphs and marked maps.

Proposition 2.3 ([2, 9]). *Let \mathbb{P} be a unimodular probability measure on \mathcal{G}_\bullet with $\mathbb{P}[\deg(\rho)] < \infty$ and let \mathbb{P}_{rev} denote the $\deg(\rho)$ -biasing of \mathbb{P} , defined by*

$$\mathbb{P}_{\text{rev}}((G, \rho) \in \mathcal{A}) := \frac{\mathbb{P}[\deg(\rho) \mathbb{1}((G, \rho) \in \mathcal{A})]}{\mathbb{P}[\deg(\rho)]}$$

for every Borel set $\mathcal{A} \subseteq \mathcal{G}_\bullet$. Then \mathbb{P}_{rev} is reversible. Conversely, if \mathbb{P} is a reversible probability measure on \mathcal{G}_\bullet , then biasing \mathbb{P} by $\deg(\rho)^{-1}$ yields a unimodular probability measure on \mathcal{G}_\bullet .

2.4 Ergodicity

A probability measure \mathbb{P} on \mathcal{G}_\bullet is said to be **ergodic** if $\mathbb{P}(\mathcal{A}) \in \{0, 1\}$ for every event $\mathcal{A} \subseteq \mathcal{G}_\bullet$ that is invariant to changing the root in the sense that

$$(G, \rho) \in \mathcal{A} \iff (G, v) \in \mathcal{A} \text{ for all } v \in V.$$

A random rooted graph is said to be ergodic if its law is ergodic. Aldous and Lyons [2, §4] proved that the ergodic unimodular probability measures on \mathcal{G}_\bullet are exactly the extreme point of the weakly closed, convex set of unimodular probability measures on. It follows by Choquet theory that every unimodular random rooted graph is a mixture of ergodic unimodular random rooted graphs, meaning that the graph may be sampled by first sampling a random ergodic unimodular probability measure on \mathcal{G}_\bullet from some distribution, and then sampling from this randomly chosen measure. In particular, to prove almost sure statements about unimodular random rooted graphs, it suffices to consider the ergodic case. Analogous statements hold for unimodular random rooted maps, marked graphs, and marked maps.

2.5 Duality

Let \mathbb{P} be a unimodular probability measure on \mathcal{M}_\bullet such that M^\dagger is locally finite \mathbb{P} -a.s., and let (M, ρ) be a random rooted map with law \mathbb{P}_{rev} . Conditional on (M, ρ) , let η be an oriented edge of M sampled uniformly at random from the set E_ρ^\rightarrow , and let $\rho^\dagger = \eta^r$. We define $\mathbb{P}_{\text{rev}}^\dagger$ to the law of the random rooted map $(M^\dagger, \rho^\dagger)$.

Proposition 2.4 (Aldous-Lyons [2, Example 9.6]). *Let \mathbb{P} be a unimodular probability measure on \mathcal{M}_\bullet such that M^\dagger is locally finite \mathbb{P} -a.s. and $\mathbb{P}[\deg(\rho)] < \infty$. Then $\mathbb{P}_{\text{rev}}^\dagger$ is a reversible probability measure on \mathcal{M}_\bullet .*

Intuitively, under the measure \mathbb{P}_{rev} , the edge η is ‘uniformly distributed’ among the edges of M . Since the map sending each edge to its dual is a bijection, it should follow that η^\dagger is uniformly distributed among the edges of M^\dagger , making the measure $\mathbb{P}_{\text{rev}}^\dagger$ reversible also.

We define \mathbb{P}^\dagger to be the unimodular measure on \mathcal{M}_\bullet obtained as the $\deg(\rho)^{-1}$ -biasing of $\mathbb{P}_{\text{rev}}^\dagger$. We refer to \mathbb{P}^\dagger as the **dual** of \mathbb{P} and say that \mathbb{P} is **self-dual** if $\mathbb{P}^\dagger = \mathbb{P}$. We write \mathbb{E}_{rev} , $\mathbb{E}_{\text{rev}}^\dagger$, and \mathbb{E}^\dagger for the associated expectation operators. We may express \mathbb{P}^\dagger directly in terms of \mathbb{P} by

$$\mathbb{P}^\dagger((M, \rho) \in \mathcal{A}) = \mathbb{E} \left[\sum_{f \perp \rho} \deg(f)^{-1} \right]^{-1} \mathbb{E} \left[\sum_{f \perp \rho} \deg(f)^{-1} \mathbb{1}((M^\dagger, f) \in \mathcal{A}) \right] \quad (2.2)$$

for every Borel set $\mathcal{A} \subseteq \mathcal{M}_\bullet$. In particular, we can calculate

$$\begin{aligned} \mathbb{E}^\dagger[\deg(\rho^\dagger)] &= \mathbb{E} \left[\sum_{f \perp \rho} \deg(f)^{-1} \right]^{-1} \mathbb{E} \left[\sum_{f \perp \rho} \deg(f)^{-1} \deg(f) \right] \\ &= \mathbb{E} \left[\sum_{f \perp \rho} \deg(f)^{-1} \right]^{-1} \mathbb{E}[\deg(\rho)] < \infty. \end{aligned} \quad (2.3)$$

The factor $\mathbb{E}[\sum_{f \perp \rho} \deg(f)^{-1}]^{-1}$ may be thought of as the ratio of the number of vertices of M to the number of faces of M .

Example 2.5. Let M be a finite map and let \mathbb{P} be the law of the unimodular random rooted map (M, ρ) where ρ is a vertex of M chosen uniformly at random. Then \mathbb{P}^\dagger is the law of the unimodular random rooted map $(M^\dagger, \rho^\dagger)$ obtained by rooting the dual M^\dagger of M at a uniformly chosen face ρ^\dagger of M .

3 Percolations and Invariant Amenability

3.1 Markings and Percolations

Let (G, ρ) be a unimodular random rooted graph with law \mathbb{P} . Given a random rooted graph (G, ρ) and a polish space \mathbb{X} , an **\mathbb{X} -marking** of (G, ρ) is a random assignment $m : E \cup V \rightarrow \mathbb{X}$ of marks to the edges and vertices of (G, ρ) (possibly defined on some larger probability space) such that the random rooted marked graph (G, ρ, m) is unimodular. A **percolation** on (G, ρ) is a $\{0, 1\}$ -marking of (G, ρ) , which we think of as a random subgraph of G consisting of the **open** edges satisfying $\omega(e) = 1$, and open vertices satisfying $\omega(v) = 1$. We assume without loss of generality that if an edge is open then so are both of its endpoints. We call ω a **bond percolation** if $\omega(v) = 1$ for every vertex $v \in V$ almost surely. The **cluster** $K_\omega(v)$ of ω at the vertex v is the connected component of v in ω . A percolation is said to be **finitary** if all of its clusters are finite almost surely. A percolation is said to be **connected** if all open vertices form a single cluster almost surely.

Let us gather here the following simple and useful technical lemmas concerning markings and percolations.

Lemma 3.1. *Let (G, ρ, m) be a unimodular random marked graph. Let ω be a percolation on (G, ρ, m) and let $m|_{K_\omega(\rho)}$ denote the restriction of m to $K_\omega(\rho)$. Then the conditional law of $(K_\omega(\rho), \rho, m|_{K_\omega(\rho)})$ given that $\omega(\rho) = 1$ is also unimodular. In particular, if ω is finitary, then ρ is uniformly distributed on its cluster.*

Proof. Let f be a mass transport, and let g be the mass transport

$$g(G, u, v, m, \omega) := \mathbb{1}(v \in K_\omega(u)) f(K_\omega(u), u, v, m|_{K_\omega(u)}).$$

Applying the mass transport principle for (G, ρ, ω) , we have

$$\begin{aligned} \mathbb{E} \sum_{v \in K_\omega(\rho)} f(K_\omega(\rho), \rho, v, m|_{K_\omega(\rho)}) &= \mathbb{E} \sum_{v \in V(G)} g(G, \rho, v, m, \omega) = \mathbb{E} \sum_{v \in V(G)} g(G, v, \rho, m, \omega) \\ &= \mathbb{E} \sum_{v \in K_\omega(\rho)} f(K_\omega(\rho), v, \rho, m|_{K_\omega(\rho)}). \end{aligned} \quad \square$$

The next two lemmas are proved by abstract probabilistic arguments. They state (roughly) that multiple markings on a graph can be combined, and that a marking on a percolation cluster can be extended to a marking on the entire graph.

Lemma 3.2 (Coupling markings I). *Let (G, ρ) be a unimodular random rooted graph and let $\{m_i : i \in I\}$ be a set of markings of (G, ρ) indexed by a countable set I with mark spaces $\{\mathbb{X}_i : i \in I\}$. Then there exists a $\prod_{i \in I} \mathbb{X}_i$ -marking m of (G, ρ) such that $(G, \rho, \pi_i \circ m)$ and (G, ρ, m_i) have the same distribution for all $i \in I$, where π_i denotes the projection of $\prod_{i \in I} \mathbb{X}_i$ onto \mathbb{X}_i for each $i \in I$.*

Proof. One such marking is given by first sampling (G, ρ) and then sampling the markings m_i independently from their conditional distributions given (G, ρ) . The details of this construction are omitted. \square

Lemma 3.3 (Coupling markings II). *Let (G, ρ) be a unimodular random rooted graph, let ω be a connected percolation on (G, ρ) , and let m be an \mathbb{X} -marking of the unimodular random rooted graph given by sampling $(K_\omega(\rho), \rho)$ conditional on $\omega(\rho) = 1$. Then there exists a \mathbb{X} -marking \hat{m} of (G, ρ) such that the laws of $(K_\omega(\rho), \rho, \hat{m}|_{K_\omega(\rho)})$ and $(K_\omega(\rho), \rho, m)$ coincide.*

Proof. One such marking is given as follows: First sample $(K_\omega(\rho), \rho)$ from its conditional distribution given $\omega(\rho) = 1$. Then, independently, sample (G, ρ) and m independently conditional on $(K_\omega(\rho), \rho)$. Extend m to the vertices and edges of G not present in ω by setting m to be some constant $x_0 \in \mathbb{X}$ on those vertices. This yields the law of a random rooted marked graph (G, ρ, ω, m) in which the root always satisfies $\omega(\rho) = 1$.

Now, for each vertex u of G , let $v(u)$ be chosen uniformly from the set of vertices of ω closest to u , independently from everything else. (In particular, $v(u) = u$ if $\omega(u) = 1$.) Note that (by the mass transport principle), the expectation of $|\{u : v(u) = \rho\}|$ is $\mathbb{P}(\rho \in \omega)^{-1}$ and so is finite. Sample the marked graph (G, ρ, ω, m) biased by $|\{u : v(u) = \rho\}|$, and let ρ' be uniform in $\{u : v(u) = \rho\}$. Then (G, ρ', ω, m) yields the desired coupling. \square

3.2 Amenability

Recall that the **(edge) Cheeger constant** of an infinite graph $G = (V, E)$ is defined to be

$$\mathbf{i}_E(G) = \inf \left\{ \frac{|\partial_E W|}{|W|} : W \subset V \text{ finite} \right\}$$

where $\partial_E W$ denotes the set of edges with exactly one end in W . The graph is said to be **amenable** if its Cheeger constant is zero and **nonamenable** if it is positive.

Nonamenability is often too strong a condition to hold for random graphs. For example, it is easily seen that every infinite cluster of a Bernoulli percolation on a nonamenable Cayley graph is almost surely amenable, since the cluster will contain atypical ‘bad regions’ that have small expansion. In light of this, Aldous and Lyons introduced the following weakened notion of nonamenability for unimodular random graphs. (Another way to overcome this issue is to use *anchored expansion*, which is less relevant in our setting. See [2] for a comparison of the two notions.) The **invariant Cheeger constant** of an ergodic unimodular random rooted graph (G, ρ) is defined to be

$$\mathbf{i}^{\text{inv}}(G, \rho) = \inf \left\{ \mathbb{E} \left[\frac{|\partial_E K_\omega(\rho)|}{|K_\omega(\rho)|} \right] : \omega \text{ a finitary percolation on } G \right\}. \quad (3.1)$$

(This is a slight abuse of notation: $\mathbf{i}^{\text{inv}}(G, \rho)$ is really a function of the *law* of (G, ρ) .) An ergodic unimodular random rooted graph (G, ρ) is said to be **invariantly amenable** if $\mathbf{i}^{\text{inv}}(G, \rho) = 0$ and **invariantly nonamenable** otherwise. More generally, we say that a unimodular random rooted graph is invariantly amenable if its ergodic decomposition is supported on invariantly amenable unimodular random rooted graphs, and invariantly nonamenable if its ergodic decomposition is supported on invariantly nonamenable unimodular random rooted graphs. This is a property of the law of (G, ρ) and not of an individual graph. Invariant amenability and nonamenability is defined similarly for ergodic unimodular random rooted maps, marked graphs and marked maps.

Let (G, ρ) be a unimodular random rooted graph and let ω be a finitary percolation on G . Let $\deg_\omega(\rho)$ denote the degree of ρ in ω and let

$$\alpha(G, \rho) = \sup \{ \mathbb{E}[\deg_\omega(\rho)] : \omega \text{ a finitary percolation on } G \}$$

An easy application of the mass transport principle [2, Lemma 8.2] shows that for any finitary percolation

$$\mathbb{E}[\deg_\omega(\rho)] = \mathbb{E} \left[\frac{\sum_{v \in K_\omega(\rho)} \deg_\omega(v)}{|K_\omega(\rho)|} \right] = \mathbb{E} \left[\frac{\sum_{v \in K_\omega(\rho)} \deg(v) - |\partial_E K_\omega(\rho)|}{|K_\omega(\rho)|} \right].$$

It follows that, if $\mathbb{E}[\deg(\rho)] < \infty$,

$$\mathbf{i}^{\text{inv}}(G, \rho) = \mathbb{E}[\deg(\rho)] - \alpha(G, \rho).$$

Similarly, if (G, ρ) is an ergodic unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] = \infty$, then $\alpha(G, \rho) < \infty$ is a sufficient condition for $\mathbf{i}^{\text{inv}}(G, \rho)$ to be positive.

3.3 Hyperfiniteness

A closely related notion to amenability is hyperfiniteness. A unimodular random rooted graph (G, ρ) is said to be **hyperfinit** if there exists a sequence of percolations $(\omega_i)_{i \geq 1}$ of (G, ρ) such that each of the percolations ω_i is finitary, $\omega_i \subseteq \omega_{i+1}$ almost surely, and $\bigcup_{i \geq 1} \omega_i = G$ almost surely. We call such a sequence a **Følner sequence** for (G, ρ) . The following is standard in the measured equivalence relations literature and was noted to carry through to the unimodular random rooted graph setting by Aldous and Lyons [2].

Theorem 3.4 ([2, Theorem 8.5]). *Let (G, ρ) be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$. Then (G, ρ) is invariantly amenable if and only if it is hyperfinite.*

We provide a short proof below for completeness.

Proof. First suppose that (G, ρ) is hyperfinite, and let $(\omega_i)_{i \geq 0}$ be a Følner sequence for (G, ρ) . By the monotone convergence theorem, $\mathbb{E}[\deg_{\omega_n}(\rho)] \rightarrow \mathbb{E}[\deg(\rho)]$, so that

$$\alpha(G, \rho) \geq \limsup_{i \rightarrow \infty} \mathbb{E}[\deg_{\omega_i}(\rho)] = \mathbb{E}[\deg(\rho)]$$

and hence (G, ρ) is invariantly amenable. Suppose conversely that (G, ρ) is invariantly amenable. For each $i \geq 1$, there exists a finitary percolation ω_i on G such that $\mathbb{E}[\deg(\rho) - \deg_{\omega_i}(\rho)] \leq 2^{-i}$. By Lemma 3.2, these can be coupled so that $(G, \rho, \langle \omega_i \rangle_{i \geq 1})$ is unimodular. For each $i \geq 1$, let $\hat{\omega}_i = \bigcap_{j \geq i} \omega_j$. Then $\hat{\omega}_i$ is an increasing sequence of finitary percolations. Furthermore, by construction,

$$\mathbb{E}[\deg(\rho) - \deg_{\hat{\omega}_i}(\rho)] \leq \sum_{j \geq i} \mathbb{E}[\deg(\rho) - \deg_{\omega_j}(\rho)] \leq 2^{-i+1}.$$

By Borel-Cantelli, ρ is in the interior of its ω_i -cluster for all sufficiently large i almost surely. It follows by unimodularity that $\bigcup \hat{\omega}_i = G$, so that $(\hat{\omega}_i)_{i \geq 0}$ is a Følner sequence for (G, ρ) . \square

Remark. Invariantly amenable unimodular random rooted graphs are always hyperfinite whether or not $\mathbb{E}[\deg(\rho)] < \infty$. However, the graph obtained by replacing each edge of the canopy tree (i.e., the Benjamini-Schramm limit of the balls in a 3-regular tree) at height n by 2^n parallel edges is hyperfinite but nonamenable (since we use edge boundaries).

Corollary 3.5 (Theorem 1, (2) implies (7)). *Let (M, ρ) be a hyperfinite ergodic random rooted map. Then (M, ρ) is a Benjamini-Schramm limit of finite planar maps.*

Proof. Let $(\omega_i)_{i \geq 1}$ be a Følner sequence for (M, ρ) . Let M_i be the finite map with underlying graph $K_{\omega_i}(\rho)$ and map structure inherited from M . Then (M_i, ρ_i) is a finite unimodular random rooted map, with the root is uniform in M_i (by Lemma 3.1). Moreover, (M_i, ρ) converges to (M, ρ) almost surely and hence also in distribution. \square

Theorem 3.6 (Theorem 1: (7) implies (2)). *Let (M, ρ) be an ergodic unimodular random planar map that is a Benjamini-Schramm limit of finite planar maps. Then (M, ρ) is hyperfinite.*

This is a corollary to the Lipton-Tarjan planar separator theorem [57]; see [4] for details. Similarly, it is possible to deduce that item (8) of Theorem 1 implies item (2) as an application of the low-genus separator theorem of Gilbert, Hutchinson, and Tarjan [36].

3.4 Ends

Recall that an infinite connected graph $G = (V, E)$ is said to be **k -ended** if k is minimal such that for every finite set W of vertices in G , the graph induced by the complement $V \setminus W$ has at most k distinct infinite connected components. In particular, an infinite tree is one-ended if it does not contain a simple bi-infinite path, and is two-ended if it contains a unique bi-infinite path.

The following proposition, due primarily to Aldous and Lyons [2], connects the number of ends of a unimodular random rooted graph to invariant amenability.

Proposition 3.7 (Ends and Amenability [2, 4]). *Let (G, ρ) be an infinite ergodic unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$. Then G has either one, two, or infinitely many ends almost surely. If G has infinitely many ends, then (G, ρ) is invariantly nonamenable. If G has two ends almost surely, then it is almost surely recurrent and (G, ρ) is invariantly amenable.*

Similarly, a connected topological space X is said to be **k -ended** if over all compact subsets K of X , the complement $X \setminus K$ has a maximum of k connected components that are not precompact.

Lemma 3.8. *Let M be a map. Then the underlying graph of M has at least as many ends as the associated surface $S = S(M)$.*

Note that M can have more ends than $S(M)$. For example the doubly infinite path has 2 ends, and the 3-regular tree is infinitely-ended, but both associated surfaces are homeomorphic to the plane and hence one-ended.

Proof. For each compact subset K of S , let W_K be the set of vertices of M that are adjacent to an edge intersecting K . Since every face of M in S is a topological disc, each non-precompact connected component of $S \setminus K$ contains infinitely many edges of M , and there are no connections in $M \setminus W_K$ between these components. Thus $M \setminus W_K$ has at least as many infinite connected components as $S \setminus K$ has non-precompact connected components. \square

Finally, let us note the following simple topological fact.

Lemma 3.9. *Let M be a map. Then the underlying graph of M has at least as many ends as M has faces of infinite degree.*

3.5 Unimodular Couplings

Let (G_1, ρ_1) and (G_2, ρ_2) be unimodular random rooted graphs. A **unimodular coupling** of (G_1, ρ_1) and (G_2, ρ_2) is a unimodular random rooted $\{0, 1\}^2$ -marked graph $(G, \rho, \omega_1, \omega_2)$ such that ω_1 and ω_2 are both connected and the law of the subgraph (ω_i, ρ) conditioned on the event that $\omega_i(\rho) = 1$ is equal to the law of (G_i, ρ_i) for each i . We say that two unimodular random rooted graphs (G_1, ρ_1) and (G_2, ρ_2) are **coupling equivalent** if they admit a unimodular coupling. (It is not difficult to show that this is an equivalence relation, arguing along similar lines to [Lemmas 3.2](#) and [3.3](#).) Coupling equivalence is closely related to the notion of two graphings generating the same equivalence relation in the theory of measured equivalence relations (see [2, Example 9.9]). Unimodular maps are defined to be coupling equivalent if the underlying graphs are. Equivalence of marked graphs and maps is defined similarly.

Example 3.10. Let (G, ρ) be a unimodular random rooted graph and let ω be a connected bond percolation on G . Then (G, ρ) and (ω, ρ) are coupling equivalent.

Example 3.11. Let (M, ρ) be a unimodular random rooted map with locally finite dual. Then (M, ρ) is coupling equivalent to its unimodular dual: both (M, ρ) and its dual can be represented as percolations of the graph formed by combining M and its dual by adding a vertex wherever a primal edge crosses a dual edge, as well as keeping the original primal and dual edges. (see [2, Example 9.6] for a similar construction). Note that the graph containing M and M^\dagger is not a map itself.

The main use of unimodular coupling will be that hyperfiniteness and strong soficity are both preserved under unimodular coupling. The corresponding statements for measured equivalence relations are due to Elek and Lippner [33].

Proposition 3.12. *Let (G_1, ρ_1) and (G_2, ρ_2) be coupling equivalent unimodular random rooted graphs. Then (G_2, ρ_2) is hyperfinite if and only if (G_1, ρ_1) is.*

Proof. It suffices to prove that if (G, ρ) is a unimodular random rooted graph and ω is a connected percolation on (G, ρ) , then (G, ρ) is hyperfinite if and only if (ω, ρ) is hyperfinite.

Suppose first that (G, ρ) is hyperfinite and let $(\tilde{\omega}_i)_{i \geq 1}$ be a Følner sequence for G . By [Lemma 3.2](#), we may couple ω and $(\omega_i)_{i \geq 1}$ so that $(G, \rho, (\omega, (\tilde{\omega}_i)_{i \geq 1}))$ is unimodular. Under such a coupling, $(\omega \cap \tilde{\omega}_i)_{i \geq 1}$ is a Følner sequence for (ω, ρ) , and consequently (ω, ρ) is hyperfinite.

Suppose conversely that (ω, ρ) is hyperfinite, and let $(\tilde{\omega}_i)_{i \geq 1}$ be a Følner sequence for (ω, ρ) . By [Lemma 3.3](#), we may assume that $(G, \rho, \omega, (\tilde{\omega}_i)_{i \geq 1})$ is unimodular. For each vertex v of G , let $w(v)$ be a chosen uniformly from the set of vertices of ω minimising the graph distance to v in G . For each i , define a subgraph $\hat{\omega}_i$ of G by

$$\hat{\omega}_i(u, v) = 1 \iff w(u) \text{ and } w(v) \text{ are connected in } \tilde{\omega}_i. \quad (3.2)$$

Then $(\hat{\omega}_i)_{i \geq 1}$ is a Følner sequence for (G, ρ) . □

Remark. It can be deduced from [Proposition 3.12](#) and [2, Theorem 8.9] that a unimodular random rooted graph with finite expected degree is hyperfinite if and only if it is coupling equivalent to \mathbb{Z} .

One useful application of [Proposition 3.12](#) is the following.

Lemma 3.13. *Let (M, ρ) be an ergodic unimodular random rooted map. The number of infinite degree faces is either 0, 1, 2, or ∞ , and if it is 1 or 2 then (M, ρ) is hyperfinite.*

Proof. By applying [Proposition 3.7](#) and [Lemma 3.9](#), it suffices to prove that if M has a non-zero but finite number of infinite degree faces almost surely, then (M, ρ) is hyperfinite. Suppose that M has finitely many infinite degree faces almost surely. Conditional on (M, ρ) , choose one infinite face, f , uniformly at random. The boundary of this face gives a coupling of M and \mathbb{Z} , except for the fact that this boundary need not be simple.

If for every vertex v of f there exists a vertex v' of f such that v is contained in a finite component of $M \setminus \{v'\}$, then inductively there exists a sequence of vertices $(v_n)_{n \geq 1}$ of M such that v_n is contained in a finite connected component of $M \setminus \{v_{n+1}\}$ for every $n \geq 1$. We deduce in this case that M is hyperfinite by taking as a Følner sequence the sequence of random subgraphs $(\tilde{\omega}_m)_{m \geq 1}$ induced by the standard monotone coupling of Bernoulli $1 - 1/m$ site percolations on M . All clusters of ω_m are almost surely finite due to the existence of the cutpoints $(v_n)_{n \geq 1}$.

Otherwise, define a percolation ω on (M, ρ) as follows. Vertices of f that are not separated from infinity by any other vertex of f are open, as are the edges of f between such vertices, and no other vertices or edges. Then ω is connected and is isomorphic to the bi-infinite line graph \mathbb{Z} . It follows that (G, ρ) is coupling equivalent to $(\mathbb{Z}, 0)$ and hence hyperfinite by [Proposition 3.12](#) \square

3.6 Vertex Extremal Length and Recurrence of Subgraphs

Vertex extremal length was introduced by He and Schramm [44] and is closely connected to circle packing. Let G be an infinite graph. For a vertex v of G , the **vertex extremal length** from v to infinity is defined to be

$$\text{VEL}_G(v, \infty) = \sup_m \frac{\inf_{\gamma: v \rightarrow \infty} m(\gamma)^2}{\|m\|^2}, \quad (3.3)$$

where the supremum is over measures m on the vertex set of G such that $\|m\|^2 = \sum m(u)^2 < \infty$, and the infimum is over paths γ from v to ∞ in G . A connected graph is said to be **VEL parabolic** if $\text{VEL}(v \rightarrow \infty) = \infty$ for some vertex v of G (and hence for every vertex), and **VEL hyperbolic** otherwise. As noted in [44], if G is VEL parabolic then there exists a vertex measure m where $\|m\| < \infty$ and $m(\gamma) = \infty$ for every path $\gamma : v \rightarrow \infty$.

The **edge extremal length** from v to infinity is defined in the same way except that the measure m is on the edges of G , and is equal to the **effective resistance** from v to infinity. In particular, G is recurrent if and only if $\mathcal{R}_{\text{eff}}(v, \infty) = \text{EEL}(v, \infty) = \infty$ [31]. See [59] for further background.

The VEL type is monotone in the sense that subgraphs of VEL parabolic graphs are also VEL parabolic. Consequently, the following two lemmas give the corresponding implications in [Theorem 1](#).

Lemma 3.14 ([44], [Theorem 1](#), (17) implies (5)). *Let G be a locally finite, connected graph. If G is VEL hyperbolic, then it is transient. If G is transient and has bounded degrees then it is VEL hyperbolic.*

For trees, as for graphs with bounded degrees, VEL parabolicity is equivalent to recurrence:

Lemma 3.15 ([Theorem 1](#), (17) implies (6)). *Let T be a tree. Then T is transient if and only if it is VEL hyperbolic.*

Proof. If T is VEL hyperbolic then it is transient by [44]. Suppose conversely that T is VEL parabolic. Then there is a vertex measure m on T with $\|m\| < \infty$ and $m(\gamma) = \infty$ for every infinite path from v . For each edge e of T , let $u(e)$ be the endpoint of e farther from v , and define an edge measure \hat{m} on T by setting $\hat{m}(e) = m(u(e))$ for every edge e of T . Then $\|\hat{m}\|^2 = \|m\|^2 - m(v)^2 < \infty$, and every simple path $\gamma : v \rightarrow \infty$ has $\hat{m}(\gamma) = m(\gamma) - m(v) = \infty$. Thus $\mathcal{R}_{\text{eff}}(v, \infty) = \infty$ and G is recurrent. \square

Benjamini and Schramm [21] used circle packing to prove the following remarkable theorem. See [4] for an alternative proof, and [39, 6] for related results establishing recurrence in certain unbounded degree cases.

Theorem 3.16 ([21]; Theorem 1, (7) implies (17)). *Let (M, ρ) be a Benjamini-Schramm limit of finite planar maps. Then M is VEL parabolic almost surely. In particular, if M has bounded degrees, then it is almost surely recurrent for simple random walk.*

We finish this section with the following theorem.

Theorem 3.17 (Benjamini, Lyons and Schramm [15]; Aldous and Lyons [2]; Theorem 1, (3) and (4) each imply (2)). *Let (G, ρ) be an invariantly nonamenable unimodular random rooted graph. Then there exists a percolation ω on G such that a.s. every connected component of ω is nonamenable, and there exists a constant M such that $\deg(v) \leq M$ for every vertex v of G such that $\omega(v) = 1$. Furthermore, the percolation ω can be taken to be a forest.*

In particular, if a unimodular random map is nonamenable then it contains nonamenable sub-trees with bounded degrees, which in turn are VEL hyperbolic. This theorem has been very useful for studying random walks on invariantly nonamenable unimodular random rooted graphs; see [4, Section 5.1]. See [4, Section 3.3.1] for a complete proof of the first part of the theorem (in which ω is not taken to be a forest).

4 Curvature

In this section, we study the average curvature of a unimodular random rooted map, and its basic properties. Essentially, the average curvature is the average Gaussian curvature per vertex when the map is embedded on a Riemannian manifold. We begin by giving a combinatorial definition of the average curvature. We then show that the average curvature is a canonical quantity associated to the random map, in the sense that any unimodular embedding the map in a Riemannian manifold (satisfying certain integrability conditions) will result in the same average curvature.

Recall that the internal angles of a regular k -gon are given by $(k - 2)\pi/k$. We define the **angle sum** at a vertex v of a map M to be

$$\theta(v) = \theta_M(v) = \sum_{f \perp v} \frac{\deg(f) - 2}{\deg(f)} \pi.$$

(Recall that this sum counts each face with the appropriate multiplicity.) This definition extends to maps with infinite faces, with the convention that $(\infty - 2)/\infty = 1$. In the case that every face of M has degree at least 3, embed M in a manifold by gluing together regular polygons (with the ∞ -gon being the upper half-space $\{x + iy \in \mathbb{C} : y > 0\}$ with edges $\{[n, n + 1] : n \in \mathbb{Z}\}$). We then have that $\theta(v)$ is the total angle of the corners at v .




Figure 4. The centre vertices have curvature $\pi/15$, 0 and $-\pi/3$ (left to right).

Of course, in general M cannot necessarily be drawn in the Euclidean plane so that all faces are regular polygons, and the angle sum at a vertex of M need not be 2π . We define the **curvature** of M at the vertex v to be the angle sum deficit

$$\kappa(v) = \kappa_M(v) = 2\pi - \theta(v).$$

See Figure 4 for examples. The **average curvature** of a unimodular random rooted map (M, ρ) , denoted $\mathbb{K}(M, \rho)$, is the expected curvature at the root:

$$\mathbb{K}(M, \rho) = \mathbb{E}[\kappa(\rho)] = 2\pi - \mathbb{E} \left[\sum_{f \perp \rho} \frac{\deg(f) - 2}{\deg(f)} \pi \right].$$

Note that if $\mathbb{E}[\deg(\rho)]$ is finite then $\mathbb{K}(M, \rho)$ is also finite.

Example 4.1 (Finite Maps). Let M be a finite map and let ρ be a vertex of M chosen uniformly at random. Then a simple calculation using Euler's formula gives that

$$\mathbb{K}(M, \rho) = \frac{1}{|V|} \sum_{v \in V} \kappa(v) = \frac{2\pi}{|V|} (2 - 2 \text{genus}(M)). \quad (4.1)$$

If S is a Riemannian surface of genus g , and $\kappa(\cdot)$ is the associated Gaussian curvature, then the Gauss-Bonnet theorem implies that $\int_S \kappa(x) dx = 2\pi(2 - 2g)$. Consequently, for an arbitrary Riemannian metric on $S(M)$ we have

$$\mathbb{K}(M, \rho) = \frac{1}{|V|} \int_S \kappa(x) dx.$$

Thus, our definition of curvature agrees with the average Gaussian curvature per vertex, and is independent of the choice of metric. This is a special case of Theorem 4.10 below.

Example 4.2 (k -angulations). If M is a k -angulation (i.e., every face of M has degree k), then

$$\theta(v) = \frac{k-2}{k} \deg(v) \pi$$

for every vertex v of M . Consequently, if (M, ρ) is a unimodular random k -angulation, then

$$\mathbb{K}(M, \rho) = \left(2 - \frac{k-2}{k} \mathbb{E}[\deg(\rho)] \right) \pi.$$

A special case corresponding to $k = \infty$, is when M is an infinite plane tree, the angle sum at a vertex v of M is simply $\theta(v) = \pi \deg(v)$, and the average curvature is $\pi(2 - \mathbb{E} \deg(\rho))$.

Example 4.3 (Curvature of the dual measure). Let \mathbb{P} be a unimodular probability measure on \mathcal{M}_\bullet such that $\mathbb{P}[\deg(\rho)] < \infty$ and M has locally finite dual \mathbb{P} -a.s., and let \mathbb{P}^\dagger be the dual measure as defined above. By (2.2),

$$\begin{aligned}\mathbb{E}^\dagger[\kappa(\rho^\dagger)] &= 2\pi - \mathbb{E}^\dagger\left[\sum_{f \perp \rho^\dagger} \frac{\deg(f) - 2}{\deg(f)}\right]\pi \\ &= 2\pi - \mathbb{E}\left[\sum_{f \perp \rho} \frac{1}{\deg(f)}\right]^{-1} \mathbb{E}\left[\sum_{f \perp \rho} \frac{1}{\deg(f)} \sum_{v \perp f} \frac{\deg(v) - 2}{\deg(v)}\right]\pi.\end{aligned}$$

Applying the mass transport principle yields that

$$\mathbb{E}^\dagger[\kappa(\rho^\dagger)] = 2\pi - \pi \mathbb{E}\left[\sum_{f \perp \rho} \frac{1}{\deg(f)}\right]^{-1} \mathbb{E}\left[\sum_{f \perp \rho} \frac{1}{\deg(f)} \sum_{v \perp f} \frac{\deg(\rho) - 2}{\deg(\rho)}\right].$$

This can be rearranged to give

$$\mathbb{E}^\dagger[\kappa(\rho^\dagger)] = \mathbb{E}\left[\sum_{f \perp \rho} \frac{1}{\deg(f)}\right]^{-1} \mathbb{E}[\kappa(\rho)]$$

Recall that $\mathbb{E}\left[\sum_{f \perp \rho} \frac{1}{\deg(f)}\right]^{-1}$ is interpreted as the ratio between vertices of M to vertices of M^\dagger . While the average curvature per vertex is changed, this is only since the density of vertices is changed. In particular, the average curvature of (M, ρ) under \mathbb{P} has the same sign as that of $(M^\dagger, \rho^\dagger)$ under \mathbb{P}^\dagger .

Example 4.4 (Self-dual maps). Let (M, ρ) be a self-dual unimodular random rooted map. Then (2.3) implies that $\mathbb{E}[\sum_{f \perp \rho} \deg(f)^{-1}] = 1$ and so

$$\begin{aligned}\mathbb{K}(M, \rho) &= 2\pi - \mathbb{E}\left[\sum_{f \perp \rho} \frac{\deg(f) - 2}{\deg(f)}\right]\pi \\ &= 2\pi - \mathbb{E}[\deg(\rho)]\pi + 2\mathbb{E}\left[\sum_{f \perp \rho} \deg(f)^{-1}\right]\pi = 4\pi - \mathbb{E}[\deg(\rho)]\pi.\end{aligned}$$

4.1 Curvature of submaps

Let M be a map with underlying graph G , and let z be a proper embedding of M into an orientable surface S . A **submap** of M is a map represented by a triple (H, S, z) , where H is a connected subgraph of G such that the restriction of z to H is a proper embedding of H into S (in particular, the faces are topological discs). If M is simply connected, then every connected subgraph of G is also a submap of M . (This may fail if M is not simply connected.)

Since the average curvature is curvature per vertex, it is not surprising that changing the map without adding or removing vertices does not change \mathbb{K} . Indeed, in the case of a finite map, deleting edges gives an embedding of the submap on the same surface, with the same total and hence same average curvature. Recall that we call a percolation ω on (M, ρ) a bond percolation if $\omega(v) = 1$ for every vertex v of M almost surely.

Proposition 4.5 (Curvature of random submaps). *Let (M, ρ) be a unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$ and let ω be a connected bond percolation on (M, ρ) that is almost surely a submap of M . Then*

$$\mathbb{K}(\omega, \rho) = \mathbb{K}(M, \rho).$$

Remark. It is straightforward to extend of Proposition 4.5 to submaps ω of (M, ρ) , that do not include every vertex. Let (N, ρ) be the unimodular random map obtained from (ω, ρ) by conditioning on the event $\rho \in \omega$, we have that

$$\mathbb{K}(N, \rho) = \mathbb{P}(\rho \in \omega)^{-1} \mathbb{K}(M, \rho).$$

Proof. First suppose that M and ω both have a locally finite duals a.s. In this case, every face of ω is a union of finitely many faces of M . For a face f of M , denote by \hat{f} the face of ω containing f .

We define a mass transport in three stages as follows. Each vertex v sends a mass of $(\deg(f) - 2)/\deg(f)$ to each face f incident to it, counting with multiplicity as usual. Each face f then sends all the mass it receives to \hat{f} . Finally, each face \hat{f} of ω redistributes the total mass it receives uniformly to its vertices (again, counted with multiplicity). This procedure yields a mass transport in which the mass sent from u to v is

$$\sum_{u \perp f \subset \hat{f} \perp v} \frac{\deg(f) - 2}{\deg(f) \deg(\hat{f})},$$

where the sum is over faces $f \subset \hat{f}$ containing u and v respectively. In particular, we have

$$\sum_{v \in V} \left| \sum_{u \perp f \subset \hat{f} \perp v} \frac{\deg(f) - 2}{\deg(f) \deg(\hat{f})} \right| \leq \sum_{u \perp f} \left| \frac{\deg(f) - 2}{\deg(f)} \right| \leq \deg(u),$$

so that this mass-transport (which can be negative in the presence of degree 1 faces) meets the integrability requirements needed to apply the signed Mass-Transport Principle.

The total mass sent from a vertex v is $\theta_M(v)/2\pi$. The total mass passing through a face f of M is $\deg(f) - 2$, while the total mass passing through a face \hat{f} of ω is $\sum_{f \subset \hat{f}} \deg(f) - 2$. It is easy to see (by induction or using Euler's formula) that the latter quantity equals $\deg(\hat{f}) - 2$, and it follows that the total mass received by a vertex v is

$$\sum_{\hat{f} \perp v} \frac{\deg(\hat{f}) - 2}{\deg(\hat{f})} = \frac{1}{2\pi} \theta_\omega(v).$$

The mass transport principle gives that $\mathbb{E}\theta_M(\rho) = \mathbb{E}\theta_\omega(\rho)$, establishing the claim in this case.

Next, suppose that M has locally finite dual but that ω does not. Every edge of M not in ω separates some face of ω in two. For each infinite face \hat{f} of ω , let $(v_i)_{i \in \mathbb{Z}}$ be the vertices at the corners of \hat{f} in counterclockwise order. These labels are well defined up to an additive constant. For each $n \geq 1$, define a percolation ω_n by letting $e \in \omega_n$ if and only if either $e \in \omega$ or else if e connects vertices v_i and v_j of some infinite face \hat{f} with $|i - j| \geq n$. Observe that, since M has a locally finite dual, ω_n has locally finite dual for every n . It follows from the first case above that $\mathbb{K}(\omega_n, \rho) = \mathbb{K}(M, \rho)$. The sequence of random variables $|\theta_{\omega_n}(\rho)|$ are bounded by $2\pi \deg(\rho)$, so that $\mathbb{K}(\omega_n, \rho) \rightarrow \mathbb{K}(\omega, \rho)$ as $n \rightarrow \infty$ by the dominated convergence theorem, completing the proof in this case.




Figure 5. The map M_1 is defined by filling in each infinite face of M with a system of arcs. This figure demonstrates this procedure applied to one of the infinite faces of \mathbb{Z} .

Finally, suppose that M^\dagger is not locally finite. We construct a unimodular map M' with locally finite dual such that M (and hence also ω) is a submap of M' with the same vertex set. By the previous case, $\mathbb{K}(M, \rho) = \mathbb{K}(M', \rho) = \mathbb{K}(\omega, \rho)$. To this end, assign to each corner $a = (e, e')$ of every infinite face of M an i.i.d. random variable U_a , uniform in $[0, 1]$. For each face f of M , add an edge between non-adjacent corners a and b , if there is some t so that $U_a, U_b \leq t$ and every other corner c of f between a and b has $U_c > t$. Here, adding an edge between corners means the edge connects the vertices of the corners, and is inserted between the two edges of the corners in the cyclic edge order at each endpoint. It is easy to see that this gives a triangulation of the infinite face of M , so that $(M')^\dagger$ is locally finite.

It is easy to see that these edges can be drawn in the face f with no crossings: If a, b, c, d are corners appearing in clockwise order in f and the edge (a, c) is added, then $U_b > U_c$ and so the edge (b, c) cannot be added. The probability that a corner is connected to another corner at distance k away along the face is $\binom{k+1}{2}^{-1}$, since their labels must be the two smallest of $k+1$ exchangeable labels. Thus, the expected number of edges added at a corner of an infinite face is $\sum 2\binom{k+1}{2}^{-1} = 2$. It follows that $\mathbb{E} \deg_{M'}(\rho) \leq \mathbb{E} 3 \deg_M(\rho) < \infty$, and we conclude by applying the previous case. \square

Proposition 4.6. \mathbb{K} is upper-semicontinuous on the set of unimodular maps with finite expected degree with respect to the local topology. Explicitly, Let (M_n, ρ_n) be a sequence of unimodular random rooted maps with $\mathbb{E}[\deg(\rho_n)] < \infty$ that converge weakly to a unimodular random rooted map (M, ρ) with $\mathbb{E}[\deg(\rho)] < \infty$. Then

$$\mathbb{K}(M, \rho) \geq \lim_{n \rightarrow \infty} \mathbb{K}(M_n, \rho_n).$$

Proof. First suppose that none of the maps M_n have any faces of degree 1. In this case, $\theta_{M_n}(v) \geq 0$ for every v and $n \geq 1$, and the claim follows from Fatou's lemma.

Otherwise, we say that a self-loop is **redundant** if it is contractible in $S(M)$, and all the vertices of M are in the same component of the complement of the loop. Note that if M is simply connected then every self-loop is contractible, and a redundant self-loop may surround other redundant self-loops but no other edge. Let \hat{M} and \hat{M}_n be obtained from

M and M_n respectively by deleting all redundant self-loops. [Proposition 4.5](#) implies that $\mathbb{K}(\hat{M}, \rho) = \mathbb{K}(M, \rho)$ and $\mathbb{K}(\hat{M}, \rho) = \mathbb{K}(\hat{M}_n, \rho_n) = \mathbb{K}(M_n, \rho_n)$ for all $n \geq 1$. Since the operation of removing redundant self-loops is continuous in the weak topology, (\hat{M}_n, ρ_n) converges weakly to (M, ρ) , and the claim follows from Fatou's Lemma as above. \square

Combining [Proposition 4.6](#) with the equation (4.1) has the following immediate corollary.

Corollary 4.7 ([Theorem 1](#), (8) implies (1)). *Let (M, ρ) be a unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$ that is a Benjamini-Schramm limit of a sequence of finite maps $(M_n)_{n \geq 1}$ such that*

$$\frac{\text{genus}(M_n)}{|V(M_n)|} \xrightarrow{n \rightarrow \infty} 0.$$

Then $\mathbb{K}(M, \rho) \geq 0$.

We will prove in [Section 5](#) that $\mathbb{K}(M, \rho) \leq 0$ for any infinite, simply connected, unimodular random rooted planar map with finite expected degree, so that Benjamini-Schramm limits of low genus finite maps have 0 average curvature.

4.2 Invariance of the curvature

In this section, we consider *unimodular embeddings* of unimodular random rooted maps. We define extend the notion of average curvature to such an embedding and show that, under certain integrability conditions, the average curvature associated to the embedding agrees with the average curvature that we defined combinatorially. This shows that the average curvature is a canonical quantity.

We define a **metric surface embedded map** (MSEM) to be a locally finite map M together with a proper embedding z of M into an oriented metric surface S . A rooted MSEM is a MSEM together with a distinguished root vertex. Two rooted MSEMs are isomorphic if they are isomorphic as rooted maps, and there is an orientation preserving isometry between the two surfaces sending one embedding to the other.

We define the local topology on the set of isomorphism classes of rooted MSEMs by a variation on the local Gromov-Hausdorff topology: Namely, we set the distance between two rooted MSEMs (M_1, ρ_1, S_1, z_1) and (M_2, ρ_2, S_2, z_2) to be e^{-r} , where r is maximal such that there is a map isomorphism ϕ from the (graph distance) ball of radius r around ρ_1 in M_1 to the ball of radius r about ρ_2 in M_2 and a correspondence R between the balls of (metric) radius r around $z_1(\rho_1)$ and $z_2(\rho_2)$ in S_1 and S_2 respectively, denoted $B_1(r)$ and $B_2(r)$, such that if (x, x') and $(y, y') \in R$ then $|d_1(x, y) - d_2(x', y')| \leq 1/r$, and such that if $\phi(e) = e'$ then R restricts to a correspondence between the points of e and of e' that are contained in the metric balls of radius r around $z_1(\rho_1)$ and $z_2(\rho_2)$ respectively.

We define doubly rooted MSEMs, the local topology on the set of isomorphism classes of doubly rooted MSEMs, and unimodular random rooted MSEMs similarly to the graph case. It is straightforward (but rather tedious) to encode the structure of a MSEM as a marking of the underlying map of the MSEM, so that all of the usual machinery of unimodularity transports to this setting. If (M, ρ) is a unimodular random rooted map and (M, S, z, ρ) is a unimodular random rooted MSEM with underlying map (M, ρ) , we call (S, z) a **unimodular embedding** of M .

In practice, unimodular embeddings often arise as measurable, automorphism equivariant functions of the map, such as the embeddings given by circle packing and the conformal

embedding. These are easily seen to be unimodular since every mass transport on the associated MSEM is induced by a mass transport on the map.

Example 4.8. Given a map M such that every face of M has degree at least three, a natural way to embed it on a surface is as follows. Associate to each face f of degree d a regular d -gon with sides of length 1. If two faces share an edge, we identify the corresponding edges of the polygons. For infinite faces, the corresponding polygon is a half-plane $\{z \in \mathbb{C} : \text{Im}(z) \geq 0\}$ with edges $\{[n, n+1] : n \in \mathbb{Z}\}$ along the boundary, which we think of as a regular ∞ -gon. In this surface, there is no curvature on any of the faces, as they are all flat pieces of \mathbb{R}^2 . The surface is also smooth along the edges, as two faces are glued along straight segments of unit length. Thus, all of the curvature of the surface is concentrated on the vertices, and in fact the atom of curvature $\kappa_s(v)$ at v (see below) is exactly the value $\kappa(v)$ that we defined combinatorially at the beginning of this section.

Example 4.9 (glued discs). Similarly to the above, we can associate to each face of M of degree d a disc of circumference d (a half-plane if $d = \infty$). Split the boundary of each disc to unit length segments, and glue discs along these segments. This is well defined also for maps with faces of degree 1 and 2. In this case, the singular curvature at a vertex is $\kappa_s(v) = \pi(2 - \deg(v))$, but there is also positive Gaussian curvature supported on the edges.

The two examples above are jointly generalized by constructions based on gluing (possibly random) shapes assigned to faces. In what follows, we restrict ourselves to surfaces that have a smooth Riemannian metric, except possibly for conical singularities at vertices of the map, and assume that the edges of the map are embedded as smooth curves in the surface. We call a unimodular embedding of a map satisfying these conditions a **smooth embedding**. For example, any construction based on gluing polygons yields a smooth embedding.

Recall a corner is a pair (e, e') of oriented edges such that e' is the edge following e in the counterclockwise order at e^- . Given a smooth embedding of a map M , for each corner of M , we let $\text{ang}(e, e')$ be the angle between e and e' at v . In a smooth embedding, every oriented edge e of the map has a well defined total geodesic curvature in the embedding, which we denote $\kappa_g(e)$.

A smooth Riemannian surface with conical singularities has a Gaussian curvature κ associated with its metric, which is a signed measure on the surface (see e.g. [3, Chapter 5]). The curvature measure κ is absolutely continuous with respect to the area measure on the surface except for possible atoms at the conical singularities at the vertices. We denote by $\kappa_s(v)$ be the atom of curvature at a vertex v , and by $\kappa(f)$ be the total curvature of the face f , which is well-defined if f has finite degree. In a smooth embedding of a map, the mass of the atom of curvature at a vertex is given by

$$\kappa_s(v) = 2\pi - \sum_{e: e^- = v} \text{ang}(e).$$

Given a smooth embedding of a map with a locally finite dual, we define the **total curvature** at a vertex v by

$$\kappa(v) = \kappa_s(v) + \sum_{f \perp v} \frac{\kappa(f)}{\deg(f)}.$$

We are now ready to state our “invariance of the average curvature” theorem.

Theorem 4.10. *Let (M, ρ) be a unimodular random rooted map with locally finite dual and let (S, z) be a smooth unimodular embedding of M . Suppose further that either*

- (1) $\kappa_g(e) = 0$ for every oriented edge e of M , and $\kappa(f) \leq 0$ for every face f of M , or
- (2) $\sum_{e: e^- = \rho} |\kappa_g(e)|$ and $|\kappa(\rho)|$ both have finite expectation.

Then

$$\mathbb{E}[\kappa(\rho)] = \mathbb{K}(M, \rho).$$

We remark that it is possible to use Theorem 4.10 to prove Proposition 4.5 by embedding M in the polygonal manifold associated to ω . We have restricted Theorem 4.10 to maps with locally finite duals mainly for the sake of brevity. It is possible to generalize this identity to maps with infinite faces under suitable integrability assumptions.

The necessity of some integrability condition is demonstrated by the following example. Consider the map obtained by a unimodular triangulation of each of the two infinite faces of \mathbb{Z} (e.g. as in Figure 5). This map admits a unimodular embedding in either the Euclidean or hyperbolic plane, by drawing the edges of \mathbb{Z} as segments of length one along a bi-infinite geodesic, and drawing the other edges of the map as semi-circular arcs. This embedding is a measurable function of the map, and is therefore unimodular. There are no conical singularities in either case, but $\sum_{f \ni v} \kappa(f)/\deg(f)$ is 0 in the Euclidean embedding and negative at every vertex in the hyperbolic embedding, violating invariance of the curvature. In fact, it is easily seen that the map has average degree 6, so that $\mathbb{K}(M, \rho) \geq 0$. It is also easy to see that the hyperbolic embedding does not satisfy the integrability requirements of Theorem 4.10.

Proof. The proof relies on a generalization of the mass transport that we used to prove Proposition 4.5 in the case that both duals were locally finite. Now, however, we need to allow for the curvature of the faces and the geodesic curvature of the edges, using the fact that the angle sum around a face is related to these quantities by the Gauss-Bonnet Theorem (see e.g. [3, Chapter VI, Section 7]).

Each corner (e, e') of M is associated with a vertex $v = e^-$ and face $f = e^r$. The Gauss-Bonnet Theorem implies that for any face f of M ,

$$\sum_{e: e^r = f} [\pi - \text{ang}(e, e')] + \kappa(f) + \sum_{e: e^r = f} \kappa_g(e) = 2\pi.$$

Here, the sum over edges is taken for the directed clockwise cycle around f . We rewrite this identity as

$$\sum_{e: e^r = f} \left[\text{ang}(e, e') - \frac{\kappa(f)}{\deg(f)} - \kappa_g(e) \right] = (\deg(f) - 2)\pi. \quad (4.2)$$

We now define two mass transports on M . First, for each directed edge $e = (u, v)$, let u send to v a mass of $\kappa_g(e)$. The expected mass sent from ρ is $\mathbb{E} \sum_{e^- = \rho} \kappa_g(e)$, while the expected mass received by ρ is $\mathbb{E} \sum_{e^- = \rho} \kappa_g(-e) = -\mathbb{E} \sum_{e^- = \rho} \kappa_g(e)$. Equating the two we find that

$$\mathbb{E} \sum_{e^- = \rho} \kappa_g(e) = 0. \quad (4.3)$$

The integrability condition needed to apply the signed Mass-Transport Principle here is given by condition (2), while the equality is trivial under condition (1).

Our second mass transport is defined in two stages as follows. For each vertex u and corner (e, e') at u with $e^r = f$, the vertex u sends to the face f a mass of $\text{ang}(e) - \frac{\kappa(f)}{\deg(f)} - \kappa_g(e)$.

Each face then redistributes the mass it receives uniformly to the vertices at its corners, with the usual multiplicity. This defines a mass transport ϕ , given explicitly by

$$\phi(u, v) = \sum_{f: f \perp u, f \perp v} \sum_{e: e^- = v, e^+ = f} \frac{1}{\deg(f)} \left[\text{ang}(e, e') - \frac{\kappa(f)}{\deg(f)} - \kappa_g(e) \right].$$

Summing the mass sent from ρ through a face f to the corners of f gives

$$\sum_v \phi(\rho, v) = \sum_{e^- = \rho} \text{ang}(e, e') - \frac{\kappa(f)}{\deg(f)} - \kappa_g(e).$$

If (1) holds then $\phi(u, v)$ is non-negative for every u and v , while if (2) holds then the expectation $\mathbb{E} \sum_{v \in V} |\phi(\rho, v)|$ is finite, and we can apply the Mass-Transport Principle to ϕ in either case. Taking expectations and using (4.3), we have

$$\mathbb{E} \sum_v \phi(\rho, v) = \mathbb{E} \left[\theta(\rho) - \sum_{e^- = \rho} \frac{\kappa(f)}{\deg(f)} \right] = 2\pi - \mathbb{E} \kappa(\rho).$$

Applying (4.2), it follows that the total mass passing through a face f is $(\deg(f) - 2)\pi$, and so the expected total mass received by ρ is

$$\mathbb{E} \sum_v \phi(v, \rho) = \mathbb{E} \sum_{f \perp \rho} \frac{(\deg(f) - 2)\pi}{\deg(f)} = 2\pi - \mathbb{E} \kappa(M, \rho),$$

completing the proof. \square

Example 4.11 (Voronoi diagrams and Delaunay triangulations). We say that a set of points in the plane is in **general position** if no four points in the set lie on any given circle or line. Given a set Z of points in general position in either the Euclidean plane or the hyperbolic plane, the **Delaunay triangulation** of Z is a simple triangulation with vertex set Z , defined so that three points u, v , and w of Z form a triangle if and only if the unique disc containing u, v , and w in its boundary does not contain any other points of Z .

Suppose that Z is an isometry-invariant, locally finite point process in either the Euclidean plane or the hyperbolic plane, and let \hat{Z} be the Palm version of Z that is conditioned to have a point at the origin. Then the Delaunay triangulation of \hat{Z} , rooted at 0 is unimodular. See [17] for a study of the Poissonian case. This triangulation, embedded with geodesic segments for edges, satisfies condition (1) of Theorem 4.10, and we deduce that, unsurprisingly, Delaunay triangulations of hyperbolic point processes are hyperbolic while Delaunay triangulations of Euclidean point processes are parabolic. (This also follows from the methods of [4].) The dual of the Delaunay triangulation is the Voronoi diagram of the point process, which can also be made unimodular as in Section 2.5. The Voronoi diagram is also hyperbolic in the hyperbolic plane and parabolic in the Euclidean plane.

5 Spanning forests

It is evident from Figure 1 that the analysis of *random spanning forests* in a map M is central to the proof of Theorem 1. The free and wired uniform spanning forests FUSF and WUSF are known to encode properties of the random walk: any graph G has the intersection property if and only if its wired uniform spanning forest is almost surely connected, while G admits

non-constant harmonic functions of finite Dirichlet energy if and only if the laws of the free and wired uniform spanning forest are distinct. The (free and wired) minimal spanning forests, on the other hand, are related to Bernoulli bond percolation: $p_c(G) < p_u(G)$ if and only if the laws of the two minimal spanning forests are distinct.

5.1 Uniform Spanning Forests

Our primary method to relate the average curvature of a map M with the various probabilistic properties listed in [Theorem 1](#) is a formula relating the average degree of the *free uniform spanning forest* of M to its average curvature. To this aim, we begin with a succinct review of uniform spanning forest. We refer the reader to [14, 59] for a comprehensive treatment.

For a finite graph G , let UST_G be the uniform measure on spanning trees of G (i.e. connected subgraphs of G containing every vertex and no cycles), which is the law of a percolation on G . There are two natural ways to define infinite volume limits of the uniform spanning tree. Let $G = (V, E)$ be an infinite, locally finite, connected graph. An **exhaustion** of G is an increasing sequence $\langle V_n \rangle_{n \geq 1}$ of finite connected subsets of V such that $\bigcup_{n \geq 1} V_n = V$. Given an exhaustion $\langle V_n \rangle_{n \geq 1}$ of G , we define G_n to be the subgraph of G induced by V_n for each $n \geq 1$. The **free uniform spanning forest** measure of G is defined as the weak limit of the uniform spanning tree measures on the graphs G_n . That is, for each finite set $S \subset E$,

$$\text{FUSF}_G(S \subset T) := \lim_{n \rightarrow \infty} \text{UST}_{G_n}(S \subset T).$$

For each $n \geq 1$, we also construct a graph G_n^* from G by identifying every vertex in $V \setminus V_n$ into a single vertex ∂_n , and deleting all of the resulting self loops from ∂_n to itself. We then define the **wired uniform spanning forest** measure of G to be the weak limit of the uniform spanning tree measures on the graphs G_n^* . That is, for each finite set $S \subset E$,

$$\text{WUSF}_G(S \subset T) := \lim_{n \rightarrow \infty} \text{UST}_{G_n^*}(S \subset T).$$

The study of uniform spanning forests was pioneered by Pemantle [63], who showed that both limits exist for any graph G and in particular are independent of the choice of exhaustion.

A crucial link between the USFs and amenability is the following result of Aldous and Lyons [2, Proposition 18.14].

Theorem 5.1 ([Theorem 1](#), (2) implies (12)). *If (G, ρ) is an invariantly amenable unimodular random rooted graph, then $\text{FUSF}_G = \text{WUSF}_G$ almost surely.*

The free and wired uniform spanning forests enjoy the following properties:

Free dominates wired: The measure FUSF_G stochastically dominates the measure WUSF_G for every graph G .

Domination and subgraphs: Let H be a connected subgraph of G . Then the FUSF of H stochastically dominates the restriction of the FUSF of G to H .

Expected degree of the WUSF: The expected degree in the WUSF of the root of any unimodular random rooted graph is 2 [2, Proposition 7.3].

Note that, since a *connected* spanning forest cannot be strictly contained in another spanning forest, the stochastic domination above has the following immediate consequence.

Proposition 5.2 (Theorem 1, (13) implies (12)). *Let G be a graph. If the wired uniform spanning forest of G is connected almost surely, then the wired and free spanning forests of G coincide.*

Let (G, ρ) be a unimodular random rooted graph and \mathfrak{F} a sample of either WUSF_G or FUSF_G . If \mathfrak{F} is interpreted as a percolation on G , then the marked graph (G, ρ, \mathfrak{F}) is also unimodular. Indeed, since the definitions of FUSF_G and WUSF_G do not depend on the choice of exhaustion, for each mass transport $f : \mathcal{G}_{\bullet\bullet}^{\{0,1\}} \rightarrow [0, \infty]$, the expectations

$$f^F(G, u, v) = \text{FUSF}_G[f(G, u, v, \mathfrak{F})] \quad \text{and} \quad f^W(G, u, v) = \text{WUSF}_G[f(G, u, v, \mathfrak{F})]$$

are also measurable. Using this observation, we deduce the mass-transport principle for (G, ρ, \mathfrak{F}) from that of (G, ρ) .

Connections to random walk and potential theory. Although the uniform spanning tree of each G_n or G_n^* is connected, the limiting random subgraph can be disconnected. Indeed, Pemantle [63] proved that WUSF and FUSF of \mathbb{Z}^d coincide for all $d \geq 1$, and are connected if and only if $d \leq 4$. A complete characterisation of the connectivity of the WUSF was given by Benjamini, Lyons, Peres and Schramm [14]. A connected, locally finite graph is said to have the **intersection property** if the traces of two independent simple random walks started from any two vertices of the graph have infinite intersection almost surely (or, equivalently, if the two traces intersect almost surely). A graph is said to have the **non-intersection property** if the traces of two independent simple random walks on the graph have finite intersection almost surely.

Theorem 5.3 ([14]: Theorem 1, equivalence of (14) and (13)). *Let G be an infinite, locally finite, connected graph and let \mathfrak{F} be a sample of WUSF_G . Then \mathfrak{F} is connected almost surely if and only if G has the intersection property. Moreover, if G has the non-intersection property, then \mathfrak{F} has infinitely many connected components almost surely.*

In general, a graph may have neither of the intersection and non-intersection properties. For example, the graph formed by connecting two disjoint copies of \mathbb{Z}^3 by a single edge between their origins does not have either property. However, it is easily seen that this is not the case for reversible random rooted graphs.

Lemma 5.4. *Let (G, ρ) be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$. Then almost surely G has either the intersection property or the non-intersection property. Consequently, the wired uniform spanning forest of a unimodular random rooted graph with finite expected degree is almost surely either connected or has infinitely many connected components.*

Proof. By biasing by the degree we may assume that (G, ρ) is reversible. We may assume also that (G, ρ) is ergodic, otherwise we take an ergodic decomposition. Let $(X_n)_{n \geq 0}$ and $(X_{-n})_{n \geq 0}$ be independent random walks on G started at ρ . Then the event that the traces of $(X_n)_{n \geq 0}$ and $(X_{-n})_{n \geq 0}$ have infinite intersection is an invariant event for the stationary sequence $(G, (X_{n+k})_{n \in \mathbb{Z}})_{k \in \mathbb{Z}}$ and therefore has probability either zero or one by ergodicity (see [2, §4] and [4, Theorem 3.1]). \square

Recall that a function $h : V \rightarrow \mathbb{R}$ defined on the vertex set of a graph $G = (V, E)$ is said to be **harmonic** if

$$h(v) = \frac{1}{\deg(v)} \sum_{u \sim v} h(u)$$

for every vertex v of G , or equivalently if $(h(X_n))_{n \geq 0}$ is a martingale when $(X_n)_{n \geq 0}$ is a random walk on G . A graph is said to be **Liouville** if it does not admit any non-constant bounded harmonic functions, and **non-Liouville** otherwise. The following well-known proposition follows from the martingale convergence theorem, see [59, Exercise 14.28].

Proposition 5.5 (Theorem 1, (14) implies (10)). *Let G be a connected graph. If G has the intersection property, then G is Liouville.*

The converse of Proposition 5.5 does not hold for general graphs. For example, \mathbb{Z}^d is Liouville for all $d \geq 1$ but has the intersection property only for $d \leq 4$. However, Benjamini, Curien and Georgakopoulos [11] proved that the converse does hold for planar graphs.

Theorem 5.6 ([11]; Theorem 1, (10) implies (14)). *Let G be a planar graph. Then G is Liouville if and only if it has the intersection property.*

The **Dirichlet energy** of a function $f : V \rightarrow \mathbb{R}$ defined on the vertex set of G is defined to be

$$\mathcal{E}(f) = \frac{1}{2} \sum_{e \in E \rightarrow} (f(e^-) - f(e^+))^2.$$

The following is classical; see [59, Exercise 9.43].

Proposition 5.7 (Theorem 1, (10) implies (11)). *Let G be a connected graph. Then the bounded harmonic functions of finite Dirichlet energy are dense in the space of harmonic functions of finite Dirichlet energy with norm $\|f\|^2 = \mathcal{E}(f)$. In particular, if G admits a non-constant harmonic function of finite Dirichlet energy, then G admits a bounded non-constant harmonic function of finite Dirichlet energy.*

Benjamini, Lyons, Peres and Schramm [14] also relate analytic properties of G to the WUSF and FUSF of G .

Theorem 5.8 ([14]; Theorem 1, equivalence of items (12) (11)). *Let G be an infinite connected graph. Then the measures FUSF_G and WUSF_G are distinct if and only if G admits non-constant harmonic functions of finite Dirichlet energy.*

In general, the FUSF is far less understood as the WUSF: no criterion for its connectivity is known, nor is it known whether the number of components of the FUSF is non-random in every graph. However, for simply connected planar maps, the FUSF is relatively well understood thanks to the following duality: Given a map M and a set $W \subset E$, let $W^\dagger := \{e^\dagger \in E^\dagger : e \notin W\}$ be the set of dual edges whose corresponding primal edges are not contained in W . Observe that if t is a spanning tree of a finite planar map M , then the dual t^\dagger is a spanning tree of M^\dagger — it is connected because t has no cycles, and has no cycles because t is connected. This observation leads to the following.

Proposition 5.9 (USF Duality [14, Theorem 12.2]). *Let M be a simply connected map with locally finite dual M^\dagger and let \mathfrak{F} be a random variable with law FUSF_M . Then \mathfrak{F}^\dagger has law WUSF_{M^\dagger} .*

In general, if M is an infinite simply connected map and \mathfrak{F} is an essential spanning forest of M (that is, a spanning forest such that every component is infinite), then \mathfrak{F}^\dagger is an essential spanning forest of M^\dagger — it is a forest because every component of \mathfrak{F} is infinite, and is essential because \mathfrak{F} has no cycles. Moreover, the forest \mathfrak{F} is connected if and only if every component

of \mathfrak{F}^\dagger is one-ended. Furthermore, a combination of results from Aldous and Lyons [2] yields the following.

Proposition 5.10 (Theorem 1, (13) implies (2)). *Let (G, ρ) be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$ and let \mathfrak{F} be a sample of WUSF_G . If \mathfrak{F} is connected almost surely, then (G, ρ) is invariantly amenable.*

Proof. By Proposition 7.3 of [2], the average degree of \mathfrak{F} is 2. By Theorem 6.2 there, this implies that \mathfrak{F} has at most two ends. Since \mathfrak{F} is connected, Theorem 8.9 of [2] implies that (G, ρ) is invariantly amenable. \square

5.2 Connectivity and degrees in the FUSF

The first main result of this section relates the average curvature and the expected degree of the FUSF in a simply connected unimodular random rooted map.

Theorem 5.11. *Let (M, ρ) be an infinite, simply connected unimodular random map. Suppose that $\mathbb{E}[\deg(\rho)] < \infty$, and let \mathfrak{F} be a sample of FUSF_M . Then*

$$\mathbb{E}[\deg_{\mathfrak{F}}(\rho)] = \frac{1}{\pi} \mathbb{E}[\theta(\rho)] = 2 - \frac{1}{\pi} \mathbb{K}(M, \rho). \quad (5.1)$$

As an easy consequence we get the following component of Theorem 1.

Corollary 5.12 (Theorem 1, equivalence of (1) and (12) and $\mathbb{K} \leq 0$). *Let (M, ρ) be an infinite, simply connected unimodular random map with $\mathbb{E}[\deg(\rho)] < \infty$. Then the average curvature satisfies $\mathbb{K}(M, \rho) \leq 0$, and equals zero if and only if $\text{FUSF}_M = \text{WUSF}_M$ almost surely.*

Proof. Let \mathfrak{F} be a sample of FUSF_M . Since the expected degree of the WUSF in any unimodular random rooted graph is two, and FUSF_M stochastically dominates WUSF_M , we have that $\mathbb{E}[\deg_{\mathfrak{F}}(\rho)] \geq 2$ and that the measures FUSF_M and WUSF_M differ almost surely if and only if this inequality is strict. We conclude by applying Theorem 5.11. \square

Theorem 5.11 follows as an immediate corollary of Proposition 4.5 and the following theorem, which is the second main result of this section. In Section 5.5, we give an alternative, duality-based proof of Theorem 5.11 that does not rely on Theorem 5.13 or Proposition 4.5, and also applies to the free minimal spanning forest.

Theorem 5.13 (Connectivity of the FUSF). *Let (M, ρ) be an infinite, simply connected unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$. Then the free uniform spanning forest of M is connected almost surely.*

Theorem 5.13 is complemented by work by the second and third authors [47], who prove the corresponding theorem for all simply connected planar maps with *bounded degrees*. Since the measure FUSF_G stochastically dominates WUSF_G for every graph G , and has no cycles, we deduce the following immediate corollary.

Corollary 5.14 (Theorem 1, equivalence of (12) and (13)). *Let (M, ρ) be an infinite, simply connected unimodular random map with $\mathbb{E}[\deg(\rho)] < \infty$. Then $\text{FUSF}_M = \text{WUSF}_M$ if and only if the wired uniform spanning forest of M is connected almost surely.*




Figure 6. The maps T_4 (left) and M_4 (right).

Having finite expected degree is necessary for this to hold. Let T_n be a binary tree of height n drawn in the plane. The Benjamini-Schramm limit of T_n as n tends to infinity is known as the canopy tree, and can be thought of as an ‘infinite binary tree viewed from a leaf’. Let M_n be the finite map obtained by drawing two copies of T_n so that one is the reflection of the other, and attaching these two copies together at their leaves (see Figure 6). Replace each edge of M_n at distance k from the leaves by 3^k parallel edges, and call the resulting map M'_n . The Benjamini-Schramm limit $(M', \rho) = \lim M'_n$ exists and is formed of two similarly thickened canopy trees, attached together at their leaves.

The random walk on M' has constant drift away from the leaves, so it is possible for two random walks to be absorbed in different halves of M' . Thus M' does not have the intersection property (nor the non-intersection property), and $\text{WUSF}_{M'}$ is not connected. However, the space of bounded harmonic functions is spanned by the probability of being absorbed at either side of M' . It follows that bounded (and hence all) harmonic functions have infinite Dirichlet energy, and that $\text{FUSF}_{M'} = \text{WUSF}_{M'}$.

Towards a proof of Theorem 5.13, observe that if M has locally finite dual almost surely, then, by Proposition 5.9, it suffices to show that every component of the WUSF of M^\dagger is one-ended almost surely. Fortunately, this is known to hold in several large classes of graphs: The following was proven by the second author [46, 45]. See also earlier works by Pemantle [63], Benjamini, Lyons, Peres, and Schramm [14], and Aldous and Lyons [2]. See [58] for further one-endedness results in the deterministic setting.

Theorem 5.15 ([46, 45]). *Let (G, ρ) be a transient unimodular random rooted graph. Then every component of the wired uniform spanning forest of G is one-ended almost surely.*

Two-ended unimodular graphs necessarily have connected and two-ended WUSF_G . However, it is an open question whether WUSF components in a one-ended, recurrent, unimodular graphs can have two ends (more than two is impossible). The final result of this section settles the planar case.

Theorem 5.16. *Let (M, ρ) be a recurrent unimodular random rooted planar map with $\mathbb{E}[\deg(\rho)] < \infty$. Then the wired uniform spanning forest of M has the same number of ends as M almost surely (which is either one or two since M is recurrent).*

5.3 Proof of Theorems 5.13 and 5.16

The proofs are closely linked and are split according to whether the dual map M^\dagger is locally finite or not.

Proofs of when M^\dagger is locally finite. Let (M, ρ) be a unimodular random rooted map with $\mathbb{P}[\deg(\rho)] < \infty$, and suppose that the dual M^\dagger is locally finite almost surely. We may assume that (M, ρ) is ergodic, otherwise we take an ergodic decomposition. Let \mathfrak{F} be a sample of FUSF_M and let \mathfrak{F}^\dagger be the dual forest. Since $\mathbb{E}[\deg(\rho)] < \infty$, the law of $(M^\dagger, \rho^\dagger)$ is absolutely continuous w.r.t. a reversible random rooted map by [Proposition 2.4](#).

If M^\dagger is almost surely transient, [Theorem 5.15](#) implies that every component of \mathfrak{F}^\dagger is one-ended almost surely, and we deduce that \mathfrak{F} is connected almost surely.

If M^\dagger is almost surely recurrent, then \mathfrak{F}^\dagger is connected almost surely. Being unimodular and recurrent, this tree has at most two ends. Since M^\dagger is locally finite, M and M^\dagger have the same number of ends. If \mathfrak{F}^\dagger was two-ended, while M is one-ended, then \mathfrak{F} would have exactly two components, contradicting [Lemma 5.4](#). If M^\dagger is two-ended, then so is the spanning tree \mathfrak{F}^\dagger , giving [Theorem 5.16](#). In either case we deduce that \mathfrak{F} is connected almost surely. \square

The remainder of this section is dedicated to the proof of [Theorems 5.13](#) and [5.16](#) in the presence of infinite faces. We begin by developing a variant of Wilson's algorithm that allows us to sample the dual of the FUSF using random walks when the dual is not locally finite.

Given a graph G and a path γ in G that is either finite or transient, i.e. visits each vertex of G at most finitely many times, the **loop-erasure** $\text{LE}(\gamma)$ is formed by erasing cycles from γ chronologically as they are created. Formally, $\text{LE}(\gamma)_i = \gamma_{t_i}$ where the times t_i are defined recursively by $t_0 = 0$ and $t_i = 1 + \max\{t \geq t_{i-1} : \gamma_t = \gamma_{t_{i-1}}\}$. If the graph has multiple edges, then $\text{LE}(\gamma)$ uses the edge γ uses between times $t_i - 1$ and t_i . The loop-erasure of a simple random walk is known as **loop-erased random walk**, was introduced by Lawler [56], and was extensively studied.

Wilson's algorithm [67] is a method of sampling a uniform spanning tree of a finite graph by joining together loop-erased random walk paths. Benjamini, Lyons, Peres and Schramm [14] introduced a variant of Wilson's algorithm for sampling the WUSF of an infinite transient graph, known as **Wilson's algorithm rooted at infinity**. Let $G = (V, E)$ be a connected, locally finite graph and let $(v_i)_{i \geq 1}$ be an enumeration of the vertex set V . We sample a sequence of forests $(\mathfrak{F}_i)_{i \geq 0}$ in G recursively as follows:

- (1) If G is finite or recurrent, fix a vertex v_0 of G and let $\mathfrak{F}_0 = \{v_0\}$. If G is transient set $\mathfrak{F}_0 = \emptyset$.
- (2) Given \mathfrak{F}_i , start a simple random walk from v_i in M , independent of everything already sampled, and stopped if and when it first hits a vertex already included in \mathfrak{F}_i .
- (3) Take the loop-erasure of this random walk path, and let \mathfrak{F}_{i+1} be the union of \mathfrak{F}_i and this loop-erased path.
- (4) Let $\mathfrak{F} = \bigcup_{i \geq 0} \mathfrak{F}_i$.

The resulting random forest has law UST_G when G is finite [67] and WUSF_G when G is infinite [14].

If M^\dagger is locally finite, then we can sample FUSF_M by using Wilson's algorithm to sample \mathfrak{F}^\dagger with law WUSF_{M^\dagger} . We introduce a variation for sampling the dual of FUSF_M , when M^\dagger is not locally finite. Let F_{fin} and F_∞ be respectively the set of finite and infinite faces of M . The following is called **Wilson's algorithm rooted at $\{\infty\} \cup F_\infty$** .

Let $(f_i)_{i \geq 1}$ be an enumeration of the finite faces of M . We sample an increasing sequence of forests $(\mathfrak{F}_i^\dagger)_{i \geq 0}$ in M^\dagger recursively as follows:

- (1') Let $\mathfrak{F}_0^\dagger = F_\infty$.

- (2') Given \mathfrak{F}_i^\dagger , start a simple random walk from f_i in M^\dagger , independent of everything we have already sampled, and stopped if and when it first hits a face already included in \mathfrak{F}_i^\dagger . In particular, the walk is stopped if it hits a face of infinite degree, so that it is well-defined.
- (3') Take the loop-erasure of this random walk path, and let $\mathfrak{F}_{i+1}^\dagger$ be the union of \mathfrak{F}_i^\dagger and this loop-erased path.
- (4') Let $\mathfrak{F}^\dagger = \bigcup_{i \geq 0} \mathfrak{F}_i^\dagger$.

Proposition 5.17. *Let $M = (V, E, \sigma)$ be a simply connected map with dual M^\dagger and let \mathfrak{F}^\dagger be a random subset of E^\dagger sampled by Wilson's algorithm rooted at $\{\infty\} \cup F_\infty$. Then $\mathfrak{F} = (\mathfrak{F}^\dagger)^\dagger$ is a sample of FUSF_M .*

Proof. Let $(V_n)_{n \geq 0}$ be an exhaustion of V such that the submap of M induced by $V \setminus V_n$, denoted M_n , does not have any finite connected components for any n . The dual of M_n may be constructed from M^\dagger by identifying every face f of M that does not have all of its vertices included in V_n into a single vertex ∂_n , and deleting all the self-loops that are created. In particular, all infinite faces of M are identified into ∂_n for every $n \geq 1$.

Note that a simple random walk $(X_n)_{n \geq 0}$ on M^\dagger , started at some finite face f and stopped if and when it hits F_∞ , is either transient or killed at a finite time almost surely. Given these observations, the rest of the proof, below, proceeds similarly to the usual proof of the veracity of Wilson's algorithm rooted at infinity [14, Theorem 5.1].

Let H be a finite set of edges of M^\dagger , and let f_1, \dots, f_l be an enumeration of the set of faces f of M that are endpoints of at least one of the edges in H . Let $((X_j^i)_{j \geq 0} : i = 1, \dots, l)$ be a collection of independent random walks in M^\dagger , where the walk $(X_j^i)_{j \geq 0}$ is started at f_i and stopped the first time that it hits an infinite face of M . Run Wilson's algorithm in M_n^\dagger , rooted at ∂_n and starting with the faces f_1, \dots, f_l in that order, using the random walks $(X_j^i)_{j \geq 0}$: For each $i \in [l]$, let τ_i^n be the first time that the random walk $(X_j^i)_{j \geq 0}$ visits the portion of the spanning tree generated up to time $i - 1$, so that

$$\text{UST}_{M_n^\dagger}(H \subset T) = \mathbb{P}\left(H \subseteq \bigcup_{i \leq l} \text{LE}((X_j^i)_{j=0}^{\tau_i^n})\right).$$

Now, similarly, run Wilson's algorithm on M^\dagger rooted at $\{\infty\} \cup F_\infty$, starting with the faces f_1, \dots, f_l in that order and using the random walks $(X_j^i)_{j \geq 0}$, and let τ_i be the first time that the random walk $(X_j^i)_{j \geq 0}$ visits the portion of the spanning tree generated up to time $i - 1$ (which might now be infinite). Since the walks $(X_j^i)_{j \geq 0}$ are finite or transient almost surely, we have that

$$\tau_i^n \rightarrow \tau_i \quad \text{and} \quad \text{LE}((X_j^i)_{j=0}^{\tau_i^n}) \rightarrow \text{LE}((X_j^i)_{j=0}^{\tau_i})$$

almost surely as $n \rightarrow \infty$. It follows that

$$\begin{aligned} \text{FUSF}_M(H \subset \mathfrak{F}^\dagger) &= \lim_{n \rightarrow \infty} \text{UST}_{M_n^\dagger}(H \subset T) = \lim_{n \rightarrow \infty} \mathbb{P}\left(H \subseteq \bigcup_{i \leq l} \text{LE}((X_j^i)_{j=0}^{\tau_i^n})\right) \\ &= \mathbb{P}\left(H \subseteq \bigcup_{i \leq l} \text{LE}((X_j^i)_{j=0}^{\tau_i})\right) \end{aligned}$$

completing the proof. \square

Proposition 5.18. *Let (M, ρ) be a simply connected, unimodular random rooted map with $\mathbb{P}[\deg(\rho)] < \infty$ and suppose that the dual M^\dagger contains a vertex of infinite degree almost surely.*

Let \mathfrak{F} be a sample of FUSF_M . Then every connected component of $\mathfrak{F}^\dagger \setminus F_\infty$ is finite almost surely, and consequently \mathfrak{F} is connected almost surely.

Proof of Proposition 5.18. The main ingredient is the observation that a random walk on M^\dagger started at a finite face will hit F_∞ in a finite time almost surely.

To see this, we first bias the law of M by $\deg(\rho)$ to get a reversible map. Next, we convert M into a (possibly disconnected) graph G' with edge set E^\dagger as follows. For each edge $e^\dagger \in E^\dagger$ with an endpoint in F_∞ , replace this endpoint with a new vertex of degree one. The vertex set of G' is $F_{\text{fin}} \cup L$, where L is the set of new degree one vertices of G' corresponding to edges of M incident to F_∞ . For each edge $e^\dagger \in E^\dagger$, let $G_{e^\dagger}^\dagger$ denote the connected component of G^\dagger containing e^\dagger . Let η be chosen uniformly from among the oriented edges emanating from ρ , and let ρ' be the tail of η^\dagger in G' . A similar argument to that used in Section 2.5 shows that the random rooted graph (G', ρ') is reversible. Every connected component of G' contains a vertex in L and consequently, by stationarity, a random walk on G' visits L infinitely often almost surely. We conclude by noting that the random walk on M^\dagger started at a vertex $f \in F_{\text{fin}}$ and stopped when it first hits F_∞ can be coupled with the random walk on G' started at the same f and stopped when it first hits L so that the two hitting times agree.

To study the FUSF , bias the reversible law of (G', ρ') by $1/\deg(\rho')$ to get a unimodular graph. Wilson's algorithm rooted at $\{\infty\} \cup F_\infty$ can be seen as generating a spanning forest \mathfrak{F}^\dagger of G' , starting with $\mathfrak{F}_0^\dagger = L$ and adding loop erased random walks. Proposition 5.17 implies that \mathfrak{F}^\dagger is distributed as the dual of the free uniform spanning forest of M . The argument above implies that every component of \mathfrak{F}^\dagger contains a unique vertex of L . Consider the mass transport where each vertex of G' sends a unit mass to the vertex of L in its component of \mathfrak{F}^\dagger . Then every vertex sends mass 1, and hence no vertex receives infinite mass. Thus, every component of \mathfrak{F}^\dagger is finite, and the result follows. \square

Proofs when M^\dagger is not locally finite. If M is simply connected then Theorems 5.13 and 5.16 follow immediately from Proposition 5.17. If M is recurrent and not simply connected, then M and \mathfrak{F} must both have two ends, giving Theorem 5.16. (Theorem 5.13 makes no claim about multiply connected maps.) \square

5.4 Percolation and minimal spanning forests

While the uniform spanning forests are related to random walks, the *minimal* spanning forests are related to bernoulli bond percolation. We recall some of the connections, and refer the reader to [59, §11] for a detailed account. **Bernoulli- p bond percolation** on G , denoted ω_p , is the random subgraph of G defined by keeping each edge of G independently with probability p and deleting the rest. The Bernoulli bond percolations $\{\omega_p\}_{p \in [0,1]}$ on a graph G may be coupled monotonically by letting $\{U(e)\}_{e \in E}$ be a collection of i.i.d. Uniform($[0,1]$) random variables indexed by the edge set of G and setting $\omega_p(e) = \mathbb{1}(U(e) \leq p)$ for every $e \in E$ and $p \in [0,1]$. The **critical probability** of G is defined by

$$p_c(G) := \inf \{p : \mathbb{P}(\omega_p \text{ has an infinite connected component}) = 1\}.$$

It is well-known [2, 62] that if (G, ρ) is an ergodic unimodular random rooted graph, then for each $p \in [0,1]$ the number of infinite connected components of ω_p for any is non-random and is in $\{0, 1, \infty\}$ almost surely. Moreover, if (G, ρ) is a unimodular random rooted graph and p is such that ω_p has a unique infinite cluster almost surely, then $\omega_{p'}$ has a unique

infinite connected component almost surely for every $p' \geq p$ [41, 61, 2]. In light of this, the **uniqueness threshold** of a graph G is defined to be

$$p_u(G) = \inf\{p : \mathbb{P}(\omega_p \text{ has a unique infinite connected component}) = 1\}.$$

Note that if (G, ρ) is an ergodic, infinite unimodular random rooted graph, then the quantities $p_c(G)$ and $p_u(G)$ are non-random. It is of interest to determine which graphs have a non-uniqueness phase for Bernoulli bond percolation. The following was proven by Aldous and Lyons [2], following work by in the transitive setting by Aizenman, Kesten, and Newman [1], Burton and Keane [26], and Gandolfi, Keane, and Newman [34].

Theorem 5.19 (Theorem 1, (2) implies (16)). *Let (G, ρ) be an invariantly amenable unimodular random rooted graph. Then ω_p has at most one infinite cluster almost surely for every $p \in [0, 1]$, and in particular $p_c = p_u$ almost surely.*

Proof. This follows by combining Corollary 6.11, Lemma 6.12, Proposition 8.8, and Corollary 8.10 of [2]. \square

A long standing conjecture of Benjamini and Schramm [22] asserts conversely that every nonamenable transitive graph has $p_c < p_u$.

Lyons, Peres and Schramm [60] related the non-uniqueness phase to minimal spanning forests. Given a finite graph G and an injective **weight** function $U : E(G) \rightarrow \mathbb{R}$ the **minimal spanning tree** of G with respect to U is defined to be the spanning tree T of G minimising the total weight $\sum_{e \in T} U(e)$. Equivalently, an edge e of G is contained in T if and only if there does not exist a simple cycle in G containing e such that e maximises $U(e)$ among the edges in this cycle. We write MST_G for the distribution on spanning trees of G obtained by letting T be the minimal spanning tree of G with respect to weights $\{U(e)\}_{e \in E}$ given by i.i.d. $\text{Uniform}([0, 1])$ random variables.

This extends to infinite graphs using exhaustions, as for uniform spanning trees. Given an exhaustion $(V_n)_{n \geq 0}$ of an infinite graph G , we define the **free and wired minimal spanning forests** as the weak limits

$$\text{FMSF}_G(S \subset T) := \lim_{n \rightarrow \infty} \text{MST}_{G_n}(S \subset T)$$

and

$$\text{WMSF}_G(S \subset T) := \lim_{n \rightarrow \infty} \text{MST}_{G_n^*}(S \subset T).$$

The limits exist and do not depend on the choice of exhaustion. If (G, ρ) is a unimodular random rooted graph and \mathfrak{F} is a sample of either WMSF_G or FMSF_G , then \mathfrak{F} is a percolation on (G, ρ) . Unlike in the uniform case, both of the minimal spanning forests may also be defined directly on the infinite graph G as follows. Let $\{U(e) : e \in E\}$ be a collection of i.i.d. $\text{Uniform}([0, 1])$ random variables indexed by the edge set of G . An edge e of G is included in free minimal spanning forest of G if and only if it is not the heaviest edge in any simple cycle in G . An edge e of G is included in the wired minimal spanning forest of G if and only if it is not the heaviest edge in any simple cycle in G or in any bi-infinite simple path (or 'cycle through infinity') in G .

Theorem 5.20 (Theorem 1, equivalence of (15) and (16)). *Let (G, ρ) be an infinite unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$. Then $p_c(G) < p_u(G)$ if and only if $\text{FMSF}_G \neq \text{WMSF}_G$ if and only if there is at most one infinite cluster in Bernoulli p -percolation on G , at every $p \in [0, 1]$.*

Proof. Lyons, Peres and Schramm [60] proved that an infinite connected graph G has $\text{FMSF}_G = \text{WMSF}_G$ if and only if for ω_p has a unique infinite cluster for Lebesgue-a.e. $p \in [0, 1]$. Combining this with monotonicity of uniqueness [41, 2, Theorem 6.7] implies that if $\text{FMSF}_G = \text{WMSF}_G$ then $p_c(G) = p_u(G)$ and hence there is at most one infinite cluster at every $p \in [0, 1]$ except possibly at $p = p_c$. If G is invariantly amenable, there is at most one infinite cluster at every p by Theorem 5.19 and so also at $p = p_c$. If G is nonamenable, then by [2] there is no infinite cluster at $p = p_c$. \square

The minimal spanning forests share several properties with their uniform cousins:

Free dominates wired: The measure FMSF_G stochastically dominates the measure WMSF_G for every graph G .

Domination and subgraphs: let H be a connected subgraph of G . Then the FMSF of H stochastically dominates the restriction of the FMSF of G to H .

Expected degree of the WMSF: The expected degree in the WMSF of root of any unimodular random rooted graph is two [2, Proposition 7.3].

Amenability and boundary conditions: If (G, ρ) is an invariantly amenable random rooted graph, then $\text{FMSF}_G = \text{WMSF}_G$ almost surely [2, Proposition 18.14].

Planar duality: If M is a simply connected map with locally finite dual M^\dagger and \mathfrak{F} is a sample of FMSF_M , then \mathfrak{F}^\dagger has law WMSF_{M^\dagger} [59, §11.5].

From the above, we deduce that if (G, ρ) a unimodular random rooted graph, the measures FMSF_G and WMSF_G coincide almost surely if and only if the expected degree of ρ in the FMSF of G is two. As for the uniform spanning forests, we relate the expected degree of the FMSF to the average curvature. The proof of Theorem 5.21 is general enough to provide an alternative proof of Theorem 5.11 that does not rely on connectivity of the forest.

Theorem 5.21 (Theorem 1, equivalence of (1) and (15)). *Let (M, ρ) be an infinite simply connected unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$ and let \mathfrak{F} be a sample of FMSF_M . Then*

$$\mathbb{E}[\deg_{\mathfrak{F}}(\rho)] = 2 - \frac{1}{\pi} \mathbb{K}(M, \rho).$$

In particular, $\text{FMSF}_M = \text{WMSF}_M$ almost surely if and only if $\mathbb{K}(M, \rho) = 0$.

The equivalence of (16) and (2) in Theorem 1 can also be proven directly as follows. Let (M, ρ) be a simply connected, invariantly nonamenable unimodular random rooted map. If (M, ρ) has locally finite dual, we deduce that $p_c(M) < p_u(M)$ by applying the following two results.

Theorem 5.22 (Benjamini, Lyons, Peres and Schramm [13, 12]; Aldous and Lyons [2]). *Let (G, ρ) be an invariantly nonamenable unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$. Then ω_{p_c} does not contain any infinite connected components almost surely.*

Theorem 5.23 (Benjamini and Schramm [20, Theorem 3.1]). *Let (M, ρ) be an invariantly nonamenable, simply connected unimodular random rooted map with locally finite dual M^\dagger and suppose that $\mathbb{E}[\deg(\rho)] < \infty$. Then ω_p has a unique infinite connected component if and only if every component of ω_p^\dagger is finite. It follows that*

$$p_u(M) = 1 - p_c(M^\dagger)$$

almost surely and that ω_{p_u} contains a unique infinite connected component almost surely.

Since at p_c there is no infinite cluster, and at p_u one exists, it follows that $p_c < p_u$. We remark that Benjamini and Schramm proved their theorem for transitive planar graphs, but their proof extends immediately to our setting. Since we provide an alternative proof via the curvature, we omit further details.

If M does not have locally finite dual, then it must have infinitely many infinite faces by [Lemma 3.13](#), so that the underlying graph of M is infinitely ended. In this case, we have that $p_u(M) = 1$ almost surely (see [Proposition 7.5](#)), while $p_c(M) < 1$ by [Theorem 5.22](#).

Thus, we have the following.

Corollary 5.24 ([Theorem 1](#), (2) implies (16)). *Let (M, ρ) be an infinite, simply connected, invariantly nonamenable unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$. Then $p_c(M) < p_u(M)$ almost surely.*

5.5 Expected degree formula

We prove [Theorem 5.21](#). Every property of the minimal spanning forests that we use also holds for the uniform spanning forests, so that we also obtain an alternative proof of [Theorem 5.11](#) that does not rely on [Theorem 5.13](#).

Proof of [Theorem 5.21](#) (and [Theorem 5.11](#)).

Locally finite dual case. Let ω be a percolation on (M, ρ) , and let $\omega^\dagger = \{e^\dagger \in E^\dagger : e \notin \omega\}$ be the dual percolation. As in [Section 2.5](#), let η be chosen uniformly at random from the set E_ρ^\rightarrow of oriented edges of M emanating from ρ , let $\rho^\dagger = \eta^r$, let \mathbb{P}^{rev} be the $\deg(\rho)$ -biasing of \mathbb{P} and let \mathbb{P}^\dagger be the $\deg(\rho^\dagger)^{-1}$ biasing of \mathbb{P}^{rev} , so that $(M^\dagger, \rho^\dagger)$ is a unimodular random rooted map under \mathbb{P}^{rev} . We write \mathbb{E}^\dagger for the expectation operator associated to \mathbb{P}^\dagger .

Lemma 5.25. *For any percolation ω we have*

$$\mathbb{E}[\deg_\omega(\rho)] = \mathbb{E}[\deg(\rho)] - \mathbb{E}\left[\sum_{f \perp \rho} \deg(f)^{-1}\right] \mathbb{E}^\dagger[\deg_{\omega^\dagger}(\rho^\dagger)].$$

Proof. Observe that, since η is uniformly distributed on E_ρ^\rightarrow conditional on (M, ρ) , we have

$$\mathbb{E}[\deg_\omega(\rho)] = \mathbb{E}[\deg(\rho) \mathbb{1}(\eta \in \omega)] = \mathbb{E}[\deg(\rho) (1 - \mathbb{1}(\eta^\dagger \in \omega^\dagger))],$$

and so

$$\mathbb{E}[\deg_\omega(\rho)] = \mathbb{E}[\deg(\rho)] (1 - \mathbb{P}^{\text{rev}}(\eta^\dagger \in \omega^\dagger)).$$

Similarly, since under the measure \mathbb{P}^\dagger and conditional on $(M^\dagger, \rho^\dagger)$, η^\dagger is uniformly distributed on $E_{\rho^\dagger}^\rightarrow$,

$$\mathbb{E}^\dagger[\deg_{\omega^\dagger}(\rho^\dagger)] = \mathbb{E}^\dagger[\deg(\rho^\dagger)] \mathbb{P}^{\text{rev}}(\eta^\dagger \in \omega^\dagger).$$

It follows that

$$\mathbb{E}[\deg_\omega(\rho)] = \mathbb{E}[\deg(\rho)] \left(1 - \frac{\mathbb{E}^\dagger[\deg(\rho^\dagger)]}{\mathbb{E}[\deg(\rho)]} \mathbb{E}^\dagger[\deg_{\omega^\dagger}(\rho^\dagger)]\right).$$

Applying the expected degree formula (2.3), we deduce that

$$\mathbb{E}[\deg_\omega(\rho)] = \mathbb{E}[\deg(\rho)] \left(1 - \frac{\mathbb{E}[\sum_{f \perp \rho} \deg(f)^{-1}]}{\mathbb{E}[\deg(\rho)]} \mathbb{E}^\dagger[\deg_{\omega^\dagger}(\rho^\dagger)]\right),$$

which rearranges to give the desired expression. \square

Let \mathfrak{F} have law FMSF_M . By [Proposition 5.9](#), the dual forest \mathfrak{F}^\dagger is distributed according to WMSF_{M^\dagger} . Since the expected degree at the root of the WMSF in any unimodular random rooted graph is 2, we have $\mathbb{E}^\dagger[\deg_{\mathfrak{F}^\dagger}(\rho^\dagger)] = 2$ and consequently, by [Lemma 5.25](#),

$$\mathbb{E}[\deg_{\mathfrak{F}}(\rho)] = \mathbb{E}[\deg(\rho)] - 2\mathbb{E}\left[\sum_{f \perp \rho} \deg(f)^{-1}\right] = \mathbb{E}\left[\sum_{f \perp \rho} \frac{\deg(f) - 2}{\deg(f)}\right] = 2 - \frac{1}{\pi}\mathbb{K}(M, \rho).$$

This completes the proof in the case that the dual of M is locally finite.

Non-locally finite dual case. We begin with an upper bound that holds in greater generality. Given a unimodular random rooted map (M, ρ) with M^\dagger not locally finite, recall the construction of the map M' from the proof of [Proposition 4.5](#), in which each infinite face of M is triangulated.

Lemma 5.26. *If ω is a percolation on an infinite unimodular random rooted map (M, ρ) that is almost surely a spanning forest of M , then*

$$\mathbb{E}[\deg_\omega(\rho)] \leq 2 - \frac{1}{\pi}\mathbb{K}(M, \rho).$$

Proof. If M has a locally finite dual, then every component of ω^\dagger is infinite. It follows from [2, Theorem 6.1] that $\mathbb{E}^\dagger[\deg_{\omega^\dagger}(\rho^\dagger)] \geq 2$, and the claim follows from [Lemma 5.25](#). If M^\dagger is not locally finite, note that ω is also a percolation on the map M' , which has the same curvature as M , and the claim follows as above. \square

For the reverse inequality we approximate M by maps with locally finite duals. Delete each edge of M' not in E independently with probability $1 - \varepsilon$, and call the resulting map M'_ε . Clearly (M'_ε, ρ) converges locally to (M, ρ) as $\varepsilon \rightarrow 0$, and it is not hard to see that M'_ε has locally finite dual a.s. for every $\varepsilon > 0$. Let \mathfrak{F}_ε be a sample of $\text{FMSF}_{M'_\varepsilon}$. Since (M'_ε, ρ) has locally finite dual, we have that

$$\mathbb{E}[\deg_{\mathfrak{F}_\varepsilon}(\rho)] = 2 - \frac{1}{\pi}\mathbb{K}(M'_\varepsilon, \rho) = 2 - \frac{1}{\pi}\mathbb{K}(M, \rho),$$

where the second equality follows from [Proposition 4.5](#).

Since the underlying graph of M is a subgraph of the underlying graph of M'_ε , the forest \mathfrak{F} stochastically dominates the restriction $\mathfrak{F}_\varepsilon \cap E$ for every $\varepsilon > 0$. Hence, by the dominated convergence theorem, we have

$$\begin{aligned} \mathbb{E}[\deg_{\mathfrak{F}}(\rho)] &\geq \mathbb{E}[\deg_{\mathfrak{F}_\varepsilon}(\rho)] - \mathbb{E}[\deg_{M'_\varepsilon}(\rho) - \deg_M(\rho)] \\ &= 2 - \frac{1}{\pi}\mathbb{K}(M, \rho) - \mathbb{E}[\deg_{M'_\varepsilon}(\rho) - \deg_M(\rho)] \xrightarrow{\varepsilon \rightarrow 0} 2 - \frac{1}{\pi}\mathbb{K}(M, \rho), \end{aligned}$$

completing the proof. \square


6 The Conformal Type

Given a map M such that every face of M has degree at least three, we may form a surface $S(M)$ (as described in [Section 2.1](#)) by gluing regular unit polygons together according to the combinatorics of M , the boundaries of these polygons becoming the edges of M embedded in $S(M)$. Recall that we consider the upper half-space $\{x + iy \in \mathbb{C} : y > 0\}$ with edges $\{[n, n+1] : n \in \mathbb{Z}\}$ to be a regular ∞ -gon. The surface $S(M)$ is endowed naturally with a conformal structure by defining an atlas as follows.

- For each face f of M , we take as a chart the identity map from the interior of the regular polygon corresponding to f to itself.
- For each edge e of M , we define an open neighbourhood of the interior of e in S by adding to e the two triangles formed by the endpoints of e and the centres of the two faces adjacent to e (if either face is infinite, we interpret this triangle to be the infinite strip starting at e and perpendicular to the boundary of the face). To define a coordinate chart on this neighbourhood, we simply place the two triangles next to each other in the plane. The reason to take only a triangle and not the entire face is that this chart is well defined even if both sides of e are incident to the same face.



- For each vertex v of M , we define an open neighbourhood of v in S similarly by intersecting the corners of the faces adjacent to v with open discs of radius $1/2$ centred at v . We define a chart on this neighbourhood by first laying the corners out sequentially around the origin (with possible overlapping), and then applying the function $z \mapsto z^{2\pi/\theta(v)}$, suitably interpreted to get an injective map into the plane. The radius is chosen so that this definition remains valid when multiple corners of the same face are located at the same vertex.



The coordinate changes are easily seen to be analytic, so that this atlas does indeed define a Riemann surface structure on $S(M)$. We denote this Riemann surface by $\mathcal{R}(M)$.

The definition can be extended to maps containing faces of degree 1 and 2 in various ways. One of these is as follows: Given a map M , let \hat{M} be obtained from M by triangulating faces of degree 1 or 2 by adding a vertex inside each face of M that has degree 1 or 2, and connecting this vertex to each corner of the face. Every face of \hat{M} has degree at least three, and we define $\mathcal{R}(M) = \mathcal{R}(\hat{M})$. The map M can be embedded in $\mathcal{R}(M)$ by restricting the natural embedding of \hat{M} into $\mathcal{R}(\hat{M})$. An alternative way to deal with general maps is to apply this triangulation procedure to all faces of finite degree, resulting in a map with only faces of degree 3 or ∞ .

We remark that there are many other natural (and inequivalent) ways to associate Riemann surfaces to maps. For example, we could associate to each face of M a disc of circumference k , with boundary split into k arcs of length one corresponding to the edges, and glue adjacent faces according to arc length along their shared edges. For concreteness we proceed with the Riemann surface defined above, though our arguments apply with minor adaptations for many other surfaces associated with M .

If M is simply connected, the uniformization theorem implies that $\mathcal{R}(M)$ is conformally equivalent to the sphere, the plane or the disc, and we call M **conformally elliptic, parabolic, or hyperbolic** accordingly. The resulting embedding of M into the sphere, plane, or disc given by uniformizing $\mathcal{R}(M)$ is unique up to Möbius transformations of the sphere, plane or disc as appropriate. This is referred to as the **conformal embedding** of M . Conformal embeddings of unimodular random planar maps are conjectured to play a key role in the theory of two-dimensional quantum gravity, see for example [27] and references therein. Elliptic maps are necessarily finite – the focus of this section is to distinguish conformally parabolic maps from conformally hyperbolic maps.

The conformal parabolicity of an infinite planar map is equivalent to the recurrence of Brownian motion on $\mathcal{R}(M)$, which can heuristically serve as a proxy for recurrence of the map itself. Gill and Rohde [37] proved that every Benjamini-Schramm limit of finite planar maps with uniformly bounded face degrees is conformally parabolic almost surely. The main result of this section generalises and, together with [Corollary 3.5](#), provides a converse to their result.

Theorem 6.1 ([Theorem 1](#), equivalence of (9) and (2) in the simply connected case.). *Let (M, ρ) be an infinite, ergodic, simply connected unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$. Then M is conformally parabolic almost surely if and only if (M, ρ) is hyperfinite.*

See [Section 7](#) for the equivalence of (9) and (2) for multiply connected maps.

Remark. Although it is natural to try to prove [Theorem 6.1](#) as an application of [Theorem 4.10](#), it is difficult to verify the required integrability condition on the geodesic curvatures of edges. Indeed, we were able to do so only under a stronger moment assumption on $\deg(\rho)$.

Proof of [Theorem 6.1](#), Conformal parabolicity implies hyperfiniteness. Let (M, ρ) be a unimodular random rooted map such that M is a.s. conformally parabolic. We claim that (M, ρ) is hyperfinite (finite expected degree is not needed for this implication). Let z be a conformal embedding of M into the plane, which is uniquely determined up to translation and scaling.

Consider the image of the vertices $z(V)$. We first claim that the convex hull C of $z(V)$ is either the full plane, a half plane, or a doubly infinite strip; the remainder of the proof will be split into cases accordingly. Indeed, since $z(V)$ is locally finite, if C is any other set then there must exist a vertex v of M such that $z(v)$ is an extreme point of C . The external angles of C at these vertices are positive, add up to at most 2π , and are independent of the choice of z . Transporting a mass α from every vertex of M to every vertex such that $z(v)$ is an extreme point of C with external angle α gives a contradiction since the mass out from each vertex is bounded while vertices such that $z(v)$ is an extreme point of C receive infinite mass. We now show that in each of the three cases, M is hyperfinite.

Case 1: C is a strip. Without loss of generality we may assume the strip is parallel to \mathbb{R} . (Otherwise apply the subsequent argument to a suitable rotation of z .) Consider the bi-infinite path on $z(V)$ visiting the vertices of M in the order given by their real part (breaking ties randomly). This defines a unimodular coupling between (M, ρ) and \mathbb{Z} , and it follows from [Proposition 3.12](#) that (M, ρ) is hyperfinite in this case.

Case 2: C is a half-plane. Without loss of generality we may assume that C is the upper half-plane. We will use the linear structure of the boundary of H to show that (M, ρ) is coupling equivalent to \mathbb{Z} , so that it will follow from [Proposition 3.12](#) that (M, ρ) is hyperfinite.

Since $z(V)$ is locally finite, there exist vertices v such that the cone $\{w : \arg(z(v) - w) \in (\pi/3, 2\pi/3)\}$ contains no other points of $z(V)$. (The choice of angles is arbitrary; any cone with finite intersection with the half-plane will do.) Call such vertices **exposed**.

It follows from ergodicity and unimodularity that a.s. there are infinitely many exposed vertices. Observe that, again using the fact that $z(V)$ is locally finite, for every bounded interval $[a, b]$, there exist at most finitely many exposed vertices v that have $\Re z(v) \in [a, b]$. Define a path whose vertices are the exposed vertices of M , arranged in order of their real part. It follows that (M, ρ) and \mathbb{Z} are coupling equivalent as claimed.

Case 3: C is the full plane. We define D to be the Delaunay tessellation with vertex set given by the set of points $z(V)$. That is, D is a map embedded in \mathbb{C} with vertex set $z(V)$. Faces of D correspond to discs (or half planes) with no vertices in their interior and at least 3 vertices on their boundary, so that each face is circumscribed in the corresponding disc. Note that the isomorphism class of (D, ρ) is independent of the choice of the conformal embedding, which is unique up to homothety, and so (M, ρ) and (D, ρ) are coupling equivalent.

Since every half plane intersects $z(V)$, the faces of D are all finite. If D itself is not locally finite, then for every infinite degree vertex v of D there exists an infinite sequence of vertices u_i and closed discs or half-planes C_i such that C_i contains both $z(v)$ and $z(u_i)$ in its boundary and no points of $z(V)$ in its interior. By taking a subsequential limit, it follows from the fact that $z(V)$ is locally finite that there must exist a half-space containing $z(v)$ in its boundary and no points of $z(V)$ in its interior, contradicting the assumption that C is the full plane.

Thus D and its dual are locally finite. In this case, it follows from the measurability of the conformal embedding that (D, ρ) is a unimodular random rooted map, and that (M, ρ) and (D, ρ) are coupling equivalent. Thus, by [Proposition 3.12](#), it suffices to prove that (D, ρ) is hyperfinite.

Define a mass transport as follows. For each vertex u and face f of D incident to u , let $\text{ang}(f, u)$ be the angle of the corner of f at u . Transport a mass of $\text{ang}(f, u)/\deg(f)$ from u to each of the vertices incident to f , including u itself. The mass sent out by each vertex u is 2π . Since each face is a polygon in the plane, the mass received is

$$\sum_{f \perp u} \frac{\text{sum of internal angles of } f}{\deg(f)} = \sum_{f \perp u} \frac{\deg(f) - 2}{\deg(f)} \pi.$$

Applying the mass-transport principle yields that the average curvature of (D, ρ) is zero. It follows from the equivalence of items (1) and (2) of [Theorem 1](#) that (D, ρ) is hyperfinite as claimed. \square

For the converse, suppose M is a conformally hyperbolic, so that we have a conformal map ϕ from $\mathcal{R}(M)$ to the hyperbolic plane \mathbb{H} , which is unique up to isometries. Our strategy is to use ϕ to give a unimodular coupling of M with another map which is known to be invariantly nonamenable. One possibility is to use (as above) the Delaunay tessellation on $\phi(V)$. The mass transport above would be used to prove that D has negative mean curvature, and hence is nonamenable. The difficulty lies in establishing the local finiteness of the resulting map and its dual. Indeed with no further assumptions the resulting map can have infinite faces, and a modification is needed.

A crucial first step is establishing that $\phi(V)$ is reasonably dense in \mathbb{H} , in the sense that there is a unimodular partition of \mathbb{H} which assigns finite mean area to every vertex. For each corner c at each vertex v of M , let U_c be the quadrilateral in the face of c with corners given




Figure 7. The covering of U_e (grey) by discs (dashed boundaries) used in the proof of Lemma 6.2. Left: the case that e^r has degree three. Right: the case that e^r has infinite degree.

by v , the centre of the face, and the midpoints of the two edges forming the corner. If the face is infinite, we take U_c to be the half-infinite strip in f with right angled corners at the mid-points of the two edges of the corner (see Figure 7). We define a partition of $\mathcal{R}(M)$ by

$$U_v = \bigcup_{c \perp v} U_c,$$

Given sets $K \subset S \subset \mathcal{R}(M)$, with S simply connected, let ϕ_S be a conformal equivalence between S and the unit disc. We write $\text{area}_{\mathbb{H}}(K, S)$ for the hyperbolic area of $\phi_S(K)$ (which is independent of the choice of ϕ_S). We also write $\text{area}_{\mathbb{H}}(K) = \text{area}_{\mathbb{H}}(K, \mathcal{R}(M))$. Recall the Schwarz-Pick Lemma: If \mathcal{R} is a Riemann surface that is conformally equivalent to the disc, S is a simply connected open subset of \mathcal{R} , and K is a Borel subset of S , then the hyperbolic area of K considered as a subset of S is greater than or equal to the hyperbolic area of K considered as a subset of \mathcal{R} . Thus $\text{area}_{\mathbb{H}}(K) \leq \text{area}_{\mathbb{H}}(K, S)$.

Lemma 6.2. *There exists a constant C such that the following holds. Let M be a conformally hyperbolic, simply connected map such that every face of M has degree at least three. Then for every vertex v ,*

$$\text{area}_{\mathbb{H}}(U_v) \leq C \deg(v).$$

We shall require the following fact: For every $0 < \varepsilon < 1$, the hyperbolic area of the set $\{z \in \mathbb{C} : |z| < 1 - \varepsilon\}$, considered as a subset of the open unit disc, is given by

$$\frac{4\pi(1 - \varepsilon)^2}{1 - (1 - \varepsilon)^2} \leq \frac{2\pi}{\varepsilon}. \quad (6.1)$$

Proof. For each corner c incident to vertex v and face f , let U_c^1, \dots, U_c^4 be the following subsets of the polygon associated to f :

- (1) U_c^1 and U_c^2 are the $1/8$ neighbourhoods of the midpoints of the edges of c .
- (2) U_c^3 is the $7/16$ neighbourhood of v .
- (3) If f has finite degree, then U_c^4 is the intersection of U_c with the disc that is centred at the centre of the polygon corresponding to f and that reaches distance $1/16$ from the boundary of the polygon. If f has infinite degree, we let U_c^4 be the part of U_c at distance at least $1/16$ from the boundary of the half-plane corresponding to f .

It is easily verified by elementary trigonometry that the U_c is contained in the union $\bigcup_{i=1}^4 U_c^i$. See Figure 7 for an illustration.

We first prove that each of U_c^1, U_c^2 , and U_c^4 have uniformly bounded hyperbolic areas. For U_c^1 , let S be the ball of radius $1/4$ around the midpoint of an edge e , so that $U_c^1 \subset S$. Recall that for an edge e , the chart at e includes the two triangles formed by the endpoints of e and the centres of the faces adjacent to e . These triangles have disjoint interiors even if the same face is on both sides of e . It is easily verified that S is always contained in the domain formed by placing the two triangles together. Since U_c^1 is the ball of radius $1/8$ around the midpoint of e , it follows by Schwarz-Pick and (6.1) that $\text{area}_{\mathbb{H}}(U_c^1) \leq \text{area}_{\mathbb{H}}(U_c^1, S) \leq C$ for some constant (namely, 4π). The corresponding claim for U_c^2 follows similarly.

We next claim that $\text{area}_{\mathbb{H}}(U_c^4) \leq 16$. If the corresponding face f has finite degree, let S be the disc circumscribed in the polygon of f , and let K be the disc that has the same centre as S and radius $1/16$ smaller than that of S . The circumference of S is less than the perimeter of f , and hence the radius ρ of S is at most $\deg(f)/2\pi$. Thus, it follows from (6.1) that

$$\text{area}_{\mathbb{H}}(K, S) \leq 16 \deg(f).$$

Since U_c^4 is a wedge of K , it follows by the symmetry of S that $\text{area}_{\mathbb{H}}(U_c^4, S) \leq 16$. By Schwarz-Pick, the same bound holds for $\text{area}_{\mathbb{H}}(U_c^4)$. If f has infinite degree, then as a subset of f , U_c^4 is isometric to the set $\{x + iy : x \in [0, 1], y \geq 1/16\}$ in the half plane representation of \mathbb{H} . This set is easily verified to have hyperbolic area 16. (One could also deduce this from the finite degree case by taking a limit.)

Finally, we claim that for every vertex v , the hyperbolic area of $U_v^3 := \bigcup_{c \perp v} U_c^3$ is at most $C \deg(v)$, where C is a universal constant. Observe that U_v^3 is simply the ball of radius $7/16$ around v in $\mathcal{R}(M)$. Let S be the ball around v of radius $1/2$. The chart provided for S maps U_v^3 to a ball of radius $(7/16)^{2\pi/\theta(v)}$, while S gets mapped to the ball of radius $(1/2)^{2\pi/\theta(v)}$ with the same centre. It follows from (6.1) that

$$\text{area}_{\mathbb{H}}(U_v^3, S) \leq \frac{2\pi}{1 - (7/8)^{2\pi/\theta(v)}} \leq C\theta(v) \leq \pi C \deg(v),$$

where C is a constant. (Verifying that such a constant exists is a simple calculus exercise.) The claim now follows from Schwarz-Pick.

The lemma follows by combining the estimates given for U_v^3 and for each of U_c^1, U_c^2 , and U_c^4 for each corner at v . \square

We remark that a simple additional argument gives the slightly stronger statement that $\text{area}_{\mathbb{H}}(U_c)$ is uniformly bounded for each corner.

Proof of Theorem 6.1, Conformal hyperbolicity implies invariant nonamenability. We may assume that every face of M has degree at least three. Otherwise, we consider the map \hat{M} used to define $\mathcal{R}(M)$, all the faces of which have degree at least 3, and which is coupling equivalent to M . Let ϕ be a conformal equivalence between $\mathcal{R}(M)$ and the hyperbolic plane.

Let Z be an independent Poisson point process of intensity 1 on the hyperbolic plane, and let D be the Delaunay triangulation associated to Z . Since D is a Poisson-Delaunay triangulation of the hyperbolic plane, D is invariantly nonamenable (see [17, 4]). To prove nonamenability of M we give a unimodular coupling of M and D , using a larger graph G , and conclude by applying Proposition 3.12.

The graph G has vertex set $V \cup Z$, and has as edges the edges of M , the edges of D , and an edge connecting each $v \in V$ to every point $z \in Z \cap \phi(U_v)$. We mark the edges of G according to which of these three types they come from. Note that the law of (G, ρ) does not depend on the choice of ϕ . It follows from Lemma 6.2 that the expected number of points in

$Z \cap \phi(U_\rho)$ is finite. It is easily verified, using the measurability of the conformal embedding, that if we sample (G, ρ) biased by $1 + |Z \cap \phi(U_\rho)|$ and then let $\hat{\rho}$ be uniform on the set $\{\rho\} \cup (Z \cap \phi(U_\rho))$, then the resulting random rooted graph $(G, \hat{\rho})$ is unimodular. Similarly, if we sample (M, ρ) and Z biased by $|Z \cap \phi(U_\rho)|$, and let ρ' be uniform on the set $Z \cap \phi(U_\rho)$, then the resulting graph (D, ρ') is unimodular. Thus, we have defined a unimodular coupling between (M, ρ) and (D, ρ') . \square

7 Multiply-connected maps

7.1 The topology of unimodular random rooted maps.

In this section we study multiply-connected unimodular random rooted maps. We begin by classifying the possible topologies of the surface associated to a unimodular random rooted map (M, ρ) . Biringer and Raimbault [23] classified the possible topologies of unimodular random rooted complete, orientable, hyperbolic surfaces. Their methods readily generalise to our setting, yielding the following theorem. In fact, the proof is slightly less technical in our setting, and we provide a quick sketch below.

Theorem 7.1 (Topology of unimodular random rooted maps). *Let (M, ρ) be an infinite unimodular random rooted map. Then the surface associated to M is almost surely homeomorphic to one of the following surfaces: the plane, the cylinder, the Cantor tree, the infinite prison window, Jacob’s ladder, or the blossoming Cantor tree. The type is determined by how many ends the surface has and whether or not it is planar.*

Here, the **cantor tree** is a ‘tree made of tubes’, and is homeomorphic to the complement of the Cantor set in the sphere, the **infinite prison window** is ‘the lattice \mathbb{Z}^2 made of tubes’, **Jacob’s ladder** is ‘an infinite ladder made of tubes’, and the **blossoming Cantor tree** is a Cantor tree with a handle attached near each bifurcation. See Figure 8 for illustrations. Be warned that homeomorphism is an extremely weak notion here. For example, ‘the lattices \mathbb{Z}^d made of tubes’ are all homeomorphic to each other for all $d \geq 2$. We leave it to the reader to verify that each of the surfaces listed in Theorem 7.1 can occur as the almost sure homeomorphism class of a unimodular random rooted map.

In [23], Biringer and Raimbault must also allow for the surfaces above to be punctured at a locally finite set of points, corresponding to isolated ends of the surface. This does not occur in our setting, as the surfaces corresponding to unimodular random rooted maps do not have isolated ends.

Sketch of proof. An **end** ξ of an infinite graph G may be defined as a function that assigns a connected component ξ_K of $G \setminus K$ to each finite set of vertices K of G , and satisfies the consistency condition that $\xi_{K'} \subseteq \xi_K$ whenever $K' \supseteq K$. The **space of ends** of G , denoted $\partial \mathcal{E}(G)$ is the topological space with the set of ends of G as its underlying set and with a basis of open sets given by sets of the form $A_{K,W} = \{\xi \text{ an end of } G : \xi_K = W\}$, where $K \subset V$ is finite and $W \subset V$ is a connected component of $G \setminus K$. Note that the basis sets are also closed, so that the space of ends is always *zero-dimensional*, that is, its topology is induced by a basis of sets that are both open and closed. The space of ends of a surface is defined similarly, replacing instances of the word ‘finite’ by ‘compact’ above, and is also zero-dimensional.

It is well-known that every unimodular random rooted graph either has one, two, or infinitely many ends, and, in the last case, the space of ends does not have any isolated points

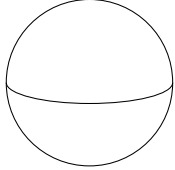
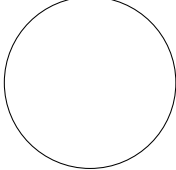
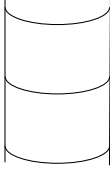
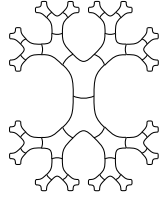
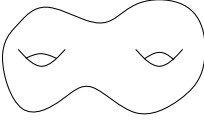
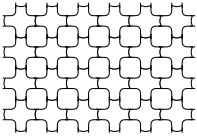

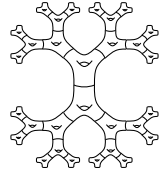
	Compact	One Ended	Two Ended	Infinitely Ended
Planar	 sphere	 disc	 cylinder	 Cantor tree
Not planar	 g -holed torus	 infinite prison window	 Jacob's ladder	 blossoming Cantor tree

Figure 8. Possible topologies of a unimodular random map. The surface $S(M)$ associated to a unimodular random map (M, ρ) is almost surely homeomorphic to one of the above.

[2, Proposition 6.10]. Since the space of ends of any graph is also compact, it follows in the last case that the space of ends is homeomorphic to the Cantor set (which, by Brouwer's Theorem [25], is the only compact, zero-dimensional Hausdorff space with no isolated points). A similar proof applies to show that if (M, ρ) is a random rooted map with associated surface $S = S(M)$, then S has either one, two, or infinitely many ends and in the last case the space of ends of S is homeomorphic to a Cantor set. Next, a standard mass transport argument (each vertex sends unit mass to the 'nearest handle') shows that if S contains handles, then the handles of S accumulate towards every end of S . That is, if S has handles then, for every compact subset K of S , every non-precompact connected component of $S \setminus K$ contains a handle.

We next apply the classification theorem for non-compact surfaces due to Kerékjártó and Richards [64], which states that if two non-compact orientable surfaces S_1 and S_2 have the same number of handles (which in our case will be zero or infinity) and there exists a homeomorphism $\phi : \partial\mathcal{E}(S_1) \rightarrow \partial\mathcal{E}(S_2)$ such that the handles of S_1 accumulate to $\xi \in \partial\mathcal{E}(S_1)$ if and only if the handles of S_2 accumulate to $\phi(\xi)$, then S_1 and S_2 are homeomorphic and ϕ extends to a homeomorphism from the ends compactification of S_1 to the ends compactification of S_2 . Thus, by the above discussion, the homeomorphism class of $S = S(M)$ is determined almost surely by its number of ends and by the existence or non-existence of handles. This yields the six different possibilities listed in the statement of the theorem (see Figure 6). \square

The main result of this section is that the average curvature of a unimodular random rooted map restricts the possible topologies of the map.

Theorem 7.2 (Topology and average curvature). *Let (M, ρ) be an ergodic unimodular random map. Then the almost sure conformal type of M is determined by its average curvature: Either*

- (1) *The average curvature of (M, ρ) is positive, in which case M is conformally elliptic and $S(M)$ is homeomorphic to the sphere almost surely,*

- (2) *the average curvature of (M, ρ) is zero, in which case M is conformally parabolic and $S(M)$ is homeomorphic to the plane, the cylinder, or the torus almost surely, or else*
- (3) *the average curvature of (M, ρ) is negative, in which case M is conformally hyperbolic $S(M)$ is homeomorphic to either the plane, the blossoming Cantor tree, Jacob's ladder, the infinite prison window, or a compact surface of genus at least two almost surely.*

The theorem will follow by combining [Theorem 6.1](#), [Theorem 7.1](#) and the notion of the universal cover of a map. Recall that a surjective, holomorphic function $\Pi : S \rightarrow S'$ between two Riemann surfaces is a **holomorphic covering** if it is locally a homeomorphism, that is, if for every $x \in S$ there exists an open neighbourhood U of S such that the restriction of Π to U is a homeomorphism between U and its image. Given a Riemann surface S , the **universal cover** of S is a simply connected Riemann surface \tilde{S} together with a covering $\Pi : \tilde{S} \rightarrow S$. The universal cover exists for any S , and is unique in the sense that if $\Pi' : \tilde{S}' \rightarrow S$ is another simply connected Riemann surface covering S , then there exists a conformal equivalence $\phi : \tilde{S}' \rightarrow \tilde{S}$ such that $\Pi' = \Pi \circ \phi$.

The universal cover of a map is defined analogously. Given a pair of maps $M = (G, \sigma)$ and $M' = (G', \sigma')$, we say that a graph homomorphism $\phi : G \rightarrow G'$ is a **map homomorphism** if $\sigma' \circ \phi = \phi \circ \sigma$, and say that ϕ is a **covering** if for every vertex v and every face f of M , the restriction of ϕ to each of $\{e \in E^- : e^- = v\}$ and $\{e \in E^+ : e^+ = f\}$ is injective. The **universal cover of a map** M is a simply connected map \tilde{M} together with a covering $\pi : \tilde{M} \rightarrow M$. Every map has a universal cover, and the universal cover of a map M is unique in the sense that if $\pi' : \tilde{M}' \rightarrow M$ is a covering from a simply connected map \tilde{M}' to M , there exists an isomorphism of maps $\phi : \tilde{M} \rightarrow \tilde{M}'$ such that $\pi' \circ \phi = \pi$.

We remark that the universal cover $\pi : \tilde{M} \rightarrow M$ of M may be constructed by lifting every edge of M in the surface $S(M)$ to the universal cover $\Pi : \tilde{S}(M) \rightarrow S(M)$ (see e.g. [42, p. 60] for the topological notion of path lifting). In particular, if $\mathcal{R}(\tilde{M})$ is the Riemann surface associated to \tilde{M} , then there exists a conformal equivalence $\Phi : \mathcal{R}(\tilde{M}) \rightarrow \tilde{\mathcal{R}}(M)$ such that $\Pi \circ \Phi \circ \tilde{z} = z \circ \pi$. (See e.g. [65, Section 9.2] for a direct construction.)

Covers of unimodular maps are unimodular. The following is proven in [4, Section 4.1].

Lemma 7.3. *Let (M, ρ) be a unimodular random rooted map. Let (\tilde{M}, π) be the universal cover of M and let $\tilde{\rho}$ be an arbitrary element of $\pi^{-1}(\rho)$. Then $(\tilde{M}, \tilde{\rho})$ is a unimodular random rooted map.*

Proof of Theorem 7.2. Observe that $\kappa_{\tilde{M}}(\tilde{\rho}) = \kappa(\rho)$ for any rooted map (M, ρ) , so that $\mathbb{K}(\tilde{M}, \tilde{\rho}) = \mathbb{K}(M, \rho)$. By applying [Theorem 6.1](#) and the classical theory of Riemann surfaces, we obtain that

- $\mathbb{K}(M, \rho) > 0$ if and only if M is finite and simply connected and $\tilde{\mathcal{R}}(M) = \mathcal{R}(M)$ is conformally equivalent to the sphere,
- $\mathbb{K}(M, \rho) = 0$ if and only if $\tilde{\mathcal{R}}(M)$ is conformally equivalent to the plane, if and only if $\mathcal{R}(M)$ is conformally equivalent to one of the plane \mathbb{C} , the cylinder \mathbb{C}/\mathbb{Z} or a torus \mathbb{C}/Λ for some lattice $\Lambda \subset \mathbb{C}$, and
- $\mathbb{K}(M, \rho) < 0$ if and only if $\tilde{\mathcal{R}}(M)$ is conformally equivalent to the hyperbolic plane (or disc).

In the last case there are many possibilities for the conformal equivalence class $\mathcal{R}(M)$; any topology other than the sphere, plane, cylinder or a torus is possible. We conclude by applying the additional topological constraints on $\mathcal{R}(M)$ imposed by [Theorem 7.1](#). \square

We next connect the topology of $S(M)$ to the number of ends of the underlying graph. If $S(M)$ is homeomorphic to the plane or the infinite prison window, it is possible for the underlying graph to have 1, 2, or infinitely many ends. (Examples in the plane case are \mathbb{Z}^2 , \mathbb{Z} and a regular tree; Add handles to these for examples with the topology of the infinite prison window.) The number of ends of the underlying graph of a map M is at least the number of ends of $S(M)$. Thus if $S(M)$ is homeomorphic to the Cantor tree or the blossoming Cantor tree, and at least two ends if $S(M)$ is homeomorphic to the cylinder or Jacob's ladder. The case of two-ended $S(M)$ is covered by the following:

Lemma 7.4. *Let (M, ρ) be a unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$. If $S(M)$ is homeomorphic to the cylinder or to Jacob's ladder almost surely, then the underlying graph of M is recurrent and two-ended almost surely.*

Proof. We bias by $\deg(\rho)$ and prove the equivalent statement for a reversible map. We may also assume that (M, ρ) is ergodic. Suppose for contradiction that M is either transient or has more than two ends almost surely. In the latter case, M has infinitely many ends almost surely, is invariantly nonamenable, and hence transient almost surely.

Since $S(M)$ is two-ended almost surely, there exists some r and D such, with positive probability, the ball $B_r(M, X_n)$ of radius r about X_n in M has degree sum at most D and the complement $S(M) \setminus z(B_r(M, X_n))$ has two non-precompact connected components, each of which is necessarily one-ended. Denote this event by A_n . By stationarity, A_n occurs for infinitely many n almost surely. Let $(n_m)_{m \geq 0}$ be a sequence of times at which A_n occurs, and such that the balls $B_r(M, X_{n_i})$ are disjoint. Let W_1, W_2 be the subgraphs of M in the two unbounded components of $S \setminus z(B_r(M, X_{n_0}))$. Since M is transient almost surely, the simple random walk $(X_n)_{n \geq 0}$ eventually stays in one of the W_i , and so W_i must be transient. However, the balls $B_r(M, X_{n_m})$ yield an infinite collection of disjoint cutsets of degree sum at most D separating ρ from infinity in the subgraph induced by W_i . Thus, this graph is recurrent by the Nash-Williams criterion [59], a contradiction. \square

7.2 [Theorem 1](#) in the multiply-connected planar case

In light of [Theorem 7.1](#) we need to consider maps where $S(M)$ is homeomorphic to either the cylinder or the Cantor tree.

Suppose that (M, ρ) is an infinite, multiply-connected unimodular random planar map with $\mathbb{E}[\deg(\rho)] < \infty$. If $\mathbb{K}(M, \rho) = 0$, then the proof of [Theorem 7.2](#) implies that $\mathcal{R}(M)$ is conformally equivalent to the cylinder. [Lemma 7.4](#) then implies that the underlying graph of M is recurrent and two-ended almost surely. We deduce that [items 2, 4–6, 10, 11, 14](#) and [17](#) of [Theorem 1](#) hold for (M, ρ) as an immediate consequence of recurrence. The remaining items of [Theorem 1](#) hold for (M, ρ) as a consequence of invariant amenability.

Now suppose that $\mathbb{K}(M, \rho) < 0$. In this case, [Theorem 7.2](#) implies that $S(M)$ is almost surely homeomorphic to the Cantor tree and consequently that the underlying graph of M is infinitely-ended almost surely by [Lemma 3.8](#), and hence that (M, ρ) is invariantly nonamenable. The following two propositions, which are well-known to experts, show that the negations of [items 11](#) and [16](#) hold for (M, ρ) . The negations of the remaining items of

[Theorem 1](#) follow from the negations of [items 11](#) and [16](#) using implications that are valid for all unimodular random rooted graphs and that we have already reviewed earlier in the paper; see the green and blue arrows in [Figure 1](#).

Proposition 7.5. *Let (G, ρ) be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$, and suppose that G is infinitely ended almost surely. Then $p_u(G) = 1$ almost surely.*

Proof of Proposition 7.5. If $p_c = 1$ the claim holds trivially, so suppose not. Fix some $p \in (p_c, 1)$, and let $I \subset V$ be the union of all infinite components of Bernoulli p -percolation. We first observe that for every finite $K \subset V$, each infinite component of $G \setminus K$ intersects I almost surely. Indeed, if not, we get a contradiction by transporting unit mass from each vertex to the nearest vertex of I . There is some finite set K so that $G \setminus K$ has multiple infinite connected components. With positive probability all edges incident to K are closed, and thus there are multiple infinite percolation components. \square

Proposition 7.6. *Let (G, ρ) be a unimodular random rooted graph with $\mathbb{E}[\deg(\rho)] < \infty$, and suppose that G is infinitely ended almost surely. Then G admits non-constant harmonic functions of finite Dirichlet energy almost surely.*

[Proposition 7.6](#) will require the following lemma.

Lemma 7.7. *Let (G, ρ) be a unimodular random rooted graph, and suppose that G is infinitely ended almost surely. Then for every finite set $K \subset V$, every infinite connected component W of $G \setminus K$ is transient.*

Proof. Since G has infinitely many ends a.s., we have that (G, ρ) is invariantly nonamenable and hence that G is transient. Let $F(u, v)$ be the probability that a simple random walk started at u hits v . Suppose for a contradiction that there exists a finite set $K \subset V$ and a recurrent infinite connected component W of $G \setminus K$. It follows that for every vertex $v \in W$, a random walk started at v must hit the set K almost surely, and hence that, for each vertex u of K , we have

$$\inf_{v \in W} F(v, u) \geq \inf_{v \in W} F(v, K) \inf_{w \in K} F(w, u) > 0.$$

Since G is transient but W is recurrent, we have that for every $\varepsilon > 0$ there exist at most finitely many vertices $v \in W$ such that $F(u, v) \geq \varepsilon$, i.e. $F(u, v) \rightarrow 0$ as $v \rightarrow \infty$ in W . Symmetry of the Green's function $\deg(x)G(x, y)$ gives

$$\frac{F(u, v)G(v, v)}{\deg(v)} = \frac{F(v, u)G(u, u)}{\deg(u)}.$$

It follows that $G(v, v)/\deg(v)$ tends to infinity as $v \rightarrow \infty$ in W .

Choose C sufficiently large that $G(\rho, \rho)/\deg(\rho) \leq C$ with positive probability. Define a mass transport by, for each vertex v of G , transporting a mass of 1 to the closest vertex to v that has $G(w, w)/\deg(w) \leq C$. (If there are multiple choices of the vertex w , choose one uniformly.) Then every vertex sends a mass of at most one but, in the situation described, some vertices receive an infinite amount of mass. This contradicts the mass transport principle. \square

Proof of Proposition 7.6. It is well-known that if G is a graph and there exists a finite set K such that $G \setminus K$ has multiple transient connected components, then G admits a non-constant harmonic Dirichlet function: The probability that the walk eventually stays in a particular connected component is such a function (see [59, Exercise 9.23]). \square

8 Soficity

We are now ready to prove [Theorem 2](#). A unimodular random rooted graph (G, ρ) is called **strongly sofic** if (G, ρ, m) is sofic for every marking m of (G, ρ) . (The constant marking shows this implies soficity.) One main ingredient to [Theorem 2](#) is the following, which was proven for Cayley graphs of free groups by [24].

Theorem 8.1 (Bowen [24]; Elek [32]; Elek and Lippner [33]; Benjamini, Lyons and Schramm [16]). *Every unimodular random rooted tree is strongly sofic.*

The following theorem is an adaptation of a related theorem of Elek and Lippner [33] in the setting of group actions. Even in our setting, it is well-known to experts that treeable unimodular random graphs (i.e., unimodular random graphs admitting a unimodular random spanning tree) are sofic.

Theorem 8.2. *Let (G_1, ρ_1) and (G_2, ρ_2) be coupling equivalent unimodular random rooted graphs. Then (G_1, ρ_1) is strongly sofic if and only if (G_2, ρ_2) is strongly sofic.*

The proof can be summarised as follows: Suppose (G_1, ρ_1) is strongly sofic, and let m be a marking of (G_2, ρ_2) . We can encode both the structure of G_2 and the marks m as a marking \hat{m} of (G_1, ρ_1) . The strong soficity of (G_1, ρ_1) allows us to approximate (G_1, ρ_1, \hat{m}) by a sequence of finite graphs. We then use this sequence to define an approximating sequence for (G_2, ρ_2, m) . Making this argument rigorous takes some care.

Proof. Since (G_1, ρ_1) and (G_2, ρ_2) can both be considered as percolations on some unimodular random graph (G, ρ) , it suffices to prove that A unimodular graph is strongly sofic if and only if a connected percolation on it is strongly sofic. Explicitly, if (G, ρ) is a unimodular random rooted graph and ω is an almost surely connected percolation on G , then (G, ρ) is strongly sofic if and only if the unimodular random rooted graph (H, ρ') obtained from (ω, ρ) by conditioning on $\omega(\rho) = 1$ is strongly sofic.

First suppose that (G, ρ) is strongly sofic and let m be a marking of (H, ρ') . By [Lemma 3.3](#), there exists a marking ω of G such that (G, ρ, ω, m) is unimodular and such that the law of (H, ρ', m) coincides with the law of (ω, ρ, m) conditional on $\omega(\rho) = 1$. Since (G, ρ) is strongly sofic, there exists a sequence of finite unimodular random marked graphs $(G_n, \rho_n, \omega_n, m_n)$ converging to (G, ρ, ω, m) in distribution. Let (H_n, ρ'_n, m_n) be the unimodular random rooted graph obtained from (ω_n, ρ_n, m_n) by conditioning on $\omega(\rho) = 1$. The sequence (H_n, ρ'_n, m_n) converges to (H, ρ', m) and, since m was arbitrary, (H, ρ') is strongly sofic.

Suppose conversely that (H, ρ') is strongly sofic and let m be an \mathbb{X} -marking of (G, ρ) . By [Lemma 3.2](#), we may assume that (G, ρ, ω, m) is unimodular. For each vertex u of G , let $v(u)$ be chosen uniformly from the set of vertices in ω that minimize the graph distance to u , and for each vertex v of ω let $U_v = \{v(u) = v\}$. Transporting mass 1 from u to $v(u)$ for every vertex u of G shows that $\mathbb{E}|U_\rho| < \infty$, and in particular U_v is finite for every vertex v of ω almost surely. Conditional on (G, ρ, ω, m) , let $\{U(v) : v \in V\} \cup \{U(e) : e \in E\}$ be a collection i.i.d. uniform $[0, 1]$ random variables indexed by the vertices and edges of G , and for each vertex v of G such that $\omega(v) = 1$, define

$$\hat{m}(v) = \{(U(u), m(u)), \{(U(e), m(e)) : e \text{ is an edge incident to } u \text{ in } G\} : u \in U_v\}.$$

If we denote the space of finite subsets of the metric space X by X^* , then \hat{m} takes values in the metric space

$$([0, 1] \times \mathbb{X} \times ([0, 1] \times \mathbb{X})^*)^*.$$

By conditioning on $\omega(\rho) = 1$, we obtain a unimodular random rooted marked graph (H, ρ, \hat{m}) . Since (H, ρ) is strongly sofic, there exists a sequence of finite unimodular random rooted marked graphs (H_n, ρ_n, \hat{m}_n) converging in distribution to (H, ρ, \hat{m}) , and we may assume that $\mathbb{E}|\hat{m}_n(\rho_n)| < \infty$ for each $n \geq 1$. For each $\varepsilon > 0$ and $R \in \mathbb{N}$, bias (H_n, ρ_n, \hat{m}_n) by $|\hat{m}(\rho)|$ and construct a finite unimodular random rooted marked graph $(G_n^{R,\varepsilon}, \rho_n, m_n^{R,\varepsilon})$ from $(H_n, \rho'_n, \hat{m}_n)$ as follows.

- (1) Let the vertex set of $G_n^{R,\varepsilon}$ be the union $\bigcup_{v \in H_n} \hat{m}(v)$, and let ρ_n be chosen uniformly from $\hat{m}(\rho'_n)$. For each vertex u of $G_n^{R,\varepsilon}$, let $u = (u^1, u^2, u^3)$ be the three coordinates of u , and let $m(u) = u^2$.
- (2) Draw an edge between two vertices u_1 and u_2 of $G_n^{R,\varepsilon}$ if and only if
 - (a) $u_1 \in \hat{m}_n(v_1)$ and $u_2 \in \hat{m}_n(v_2)$ for some vertices v_1 and v_2 of H_n that are at distance at most R in H_n , and
 - (b) there exists a pair of points

$$(t, x) \in u^3 \subset [0, 1] \times \mathbb{X} \text{ and } (s, y) \in u^3 \subset [0, 1] \times \mathbb{X}$$

such that the distance between (t, x) and (s, y) is less than ε in the product metric. If there are multiple such pairs, we draw multiple edges as appropriate.

- (3) Let $\{Z_n(e)\}$ be a collection of i.i.d. Bernoulli-1/2 random variables indexed by the edges of $G_n^{R,\varepsilon}$. For each edge e of $G_n^{R,\varepsilon}$, let (t, x) and (s, y) be the matching pair of points in $[0, 1] \times \mathbb{X}$ that led us to draw e in step (2), and let $m_n^{R,\varepsilon}(e) = x$ if $Z_n(e) = 0$ and $m_n^{R,\varepsilon}(e) = y$ if $Z_n(e) = 1$.

This construction is continuous for the local topology, and hence for each fixed ε and R , the finite marked graphs $(G_n^{R,\varepsilon}, \rho_n, m_n^{R,\varepsilon})$ converge to the marked graph $(G^{R,\varepsilon}, \rho, m^{R,\varepsilon})$ defined by applying the same procedure to (H, ρ', \hat{m}) . Taking $\varepsilon \rightarrow 0$, we obtain the marked graph (G^R, ρ, m^R) which consists of those edges of (G, ρ) whose endpoints are of distance at most R in ω . Finally, taking $R \rightarrow \infty$ we recover (G, ρ, m) . Thus, (G, ρ, m) is a weak limit of sofic unimodular random rooted marked graphs, and it follows that (G, ρ, m) is sofic. \square

Proof of Theorem 2. Let (M, ρ) be a simply connected unimodular random rooted map, let (G, ρ) be the underlying graph of M , and let \mathfrak{F} be a sample of FUSF_M . By Theorem 5.13, \mathfrak{F} is connected almost surely, and so (G, ρ) is coupling equivalent to the unimodular random rooted tree (\mathfrak{F}, ρ) . It follows from Theorems 8.1 and 8.2 that (G, ρ) is strongly sofic. By encoding the map (M, ρ) as a marking of (G, ρ) as in Section 2.2 and [2, Example 9.6], we conclude that the unimodular random rooted map (M, ρ) is also sofic. \square

9 Open Problems

We expect that the dichotomy of Theorem 1 extends to many further properties of planar unimodular random rooted maps. In this section, we discuss several such properties that might be addressed.

9.1 Rates of escape of the random walk

Can the type of a unimodular random planar map be determined by the rate of escape of the random walk? The work of Ding, Lee, and Peres [30] (together with the characterization of parabolic unimodular random planar maps as Benjamini-Schramm limits of finite planar

maps) implies that the random walk is at most diffusive on any parabolic unimodular random planar map of finite expected degree.

Theorem 9.1 ([30]). *There exists a universal constant C such that for every parabolic unimodular random rooted map (M, ρ) with $\mathbb{E}[\deg(\rho)] < \infty$, we have*

$$\mathbb{E} [\deg(\rho)d(\rho, X_n)^2] \leq Cn \mathbb{E} [\deg(\rho)]$$

for all $n \geq 0$.

On the other hand, if (M, ρ) is a hyperbolic unimodular random rooted map with $\mathbb{E}[\deg(\rho)] < \infty$ that has at most exponential growth, meaning that

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log |B(\rho, n)| < \infty,$$

then the random walk on M has positive speed, that is,

$$\lim_{n \rightarrow \infty} \frac{1}{n} d(\rho, X_n) > 0$$

a.s., where the limit exists a.s. by Kingman's subadditive ergodic theorem. (Note that the exponential growth condition always holds for graphs of bounded degree.) This can be seen in several ways: it is an easy consequence of a theorem of Benjamini, Lyons, and Schramm [15, Theorem 3.2] (see also [2, Theorem 8.13] and [4, Theorem 3.2]) that every invariantly nonamenable unimodular random rooted graph with finite expected degree and at most exponential growth has positive speed. Meanwhile, it is a result of Benjamini and Curien [9], generalizing the work of Kaimanovich, Vershik, and others [53, 52, 51, 50, 49, 48], that every non-Liouville unimodular random rooted graph with finite expected degree and at most exponential growth has positive speed.

In general, however, there do exist invariantly nonamenable, non-Liouville, unimodular random rooted graphs with finite expected degree such that the random walk has zero speed almost surely. An example of such a graph will appear in a forthcoming paper by the second author. We do not know of a planar example, which motivates the following question.

Question 9.2. *Let (M, ρ) be a hyperbolic unimodular random rooted planar map with $\mathbb{E}[\deg(\rho)] < \infty$. Does the random walk on M have positive speed almost surely?*

See [4] for a related result concerning the positivity of the speed in the hyperbolic metric induced by the circle packing of a hyperbolic unimodular random rooted triangulation with $\mathbb{E}[\deg(\rho)^2] < \infty$.

9.2 Positive harmonic functions

Theorem 1 states that the existence of non-constant bounded harmonic functions and of non-constant harmonic Dirichlet functions are both determined by the type. We conjecture that a similar result holds for positive harmonic functions.

Conjecture 9.3. *Let (M, ρ) be a parabolic unimodular random rooted map and suppose that $\mathbb{E}[\deg(\rho)] < \infty$. Then M does not admit any non-constant positive harmonic functions almost surely.*

This conjecture would follow from a positive answer to the following question.

Question 9.4. *Let (M, ρ) be a parabolic unimodular random rooted map and suppose that $\mathbb{E}[\deg(\rho)] < \infty$. Let $\langle X_n \rangle_{n \geq 0}$ be a simple random walk on M . Is every component of the complement of the trace of $\langle X_n \rangle_{n \geq 0}$ finite almost surely?*

Note that the answer to [Question 9.4](#) is trivially positive if M is recurrent.

9.3 Heat kernel decay

Let (M, ρ) be a unimodular random rooted planar map, and let $p_n(\rho, \rho)$ be the probability, given (M, ρ) , that a simple random walk on M started at ρ is at ρ at time n . To what extent are the asymptotics of $p_n(\rho, \rho)$ determined by the type of (M, ρ) ? Most hyperbolic maps encountered in practice have $p_n(\rho, \rho)$ decaying superpolynomially (when the maps have bounded degree and anchored expansion, this follows from the results of [66]). However, for each $\varepsilon > 0$, it is possible to construct a hyperbolic unimodular random rooted tree of bounded degree such that $p_n(\rho, \rho) \geq Cn^{-2-\varepsilon}$ for some $C < \infty$ and every $n \geq 1$. On the other hand, [Theorem 1](#) implies that any parabolic unimodular random planar map (M, ρ) with $\mathbb{E}[\deg(\rho)] < \infty$ must have $\mathbb{E}[\sum_{n \geq 0} np_n(\rho, \rho) \deg(\rho)] / \mathbb{E}[\deg(\rho)] = \infty$, since the left hand quantity is exactly the expected number of pairs $i < 0 < j$ such that $X_i = X_j$ when $\langle X_i \rangle_{i \in \mathbb{Z}}$ is the bi-infinite random walk once we have biased (M, ρ) to be reversible. This leads to the following question.

Question 9.5. *Let (M, ρ) be an ergodic unimodular random rooted planar map and suppose that $\mathbb{E}[\deg(\rho)] < \infty$. Do we have that*

$$\sum_{n \geq 0} np_n(\rho, \rho) = \infty$$

almost surely if and only if (M, ρ) is parabolic?

A related property that may be determined by the type is the recurrence of the critical branching random walk conditioned to survive. See e.g. [10] for background.

Question 9.6. *Let (M, ρ) be a unimodular random rooted planar map and suppose that $\mathbb{E}[\deg(\rho)] < \infty$. Is (M, ρ) parabolic if and only if the critical branching random walk on M conditioned to survive is recurrent almost surely?*

Acknowledgement

OA is supported by NSERC and the Simons Foundation. TH is supported by a Microsoft Research PhD Fellowship. AN is supported by ISF grant 1207/15, and ERC starting grant 676970. GR is supported in part by EPSRC grant EP/I03372X/1. Part of this work was conducted at the Isaac Newton Institute in Cambridge, during the programme ‘Random Geometry’ supported by EPSRC Grant Number EP/K032208/1.

References

- [1] M. Aizenman, H. Kesten, and C. M. Newman. Uniqueness of the infinite cluster and continuity of connectivity functions for short and long range percolation. *Comm. Math. Phys.*, 111(4):505–531, 1987.
- [2] D. Aldous and R. Lyons. Processes on unimodular random networks. *Electron. J. Probab.*, 12:no. 54, 1454–1508, 2007.

- [3] A. D. Aleksandrov and V. A. Zalgaller. *Intrinsic geometry of surfaces*. Translated from the Russian by J. M. Danskin. Translations of Mathematical Monographs, Vol. 15. American Mathematical Society, Providence, R.I., 1967.
- [4] O. Angel, T. Hutchcroft, A. Nachmias, and G. Ray. Unimodular hyperbolic triangulations: Circle packing and random walk. *Inventiones Mathematicae*, to appear, 2015. arXiv:1501.04677.
- [5] O. Angel and G. Ray. Classification of half planar maps. *Ann. Probab.*, to appear., 2013. arXiv:1303.6582.
- [6] O. Angel and G. Ray. The half plane uipt is recurrent. *arXiv preprint arXiv:1601.00410*, 2016.
- [7] O. Angel and O. Schramm. Uniform infinite planar triangulations. *Comm. Math. Phys.*, 241(2-3):191–213, 2003.
- [8] A. F. Beardon and K. Stephenson. Circle packings in different geometries. *Tohoku Mathematical Journal, Second Series*, 43(1):27–36, 1991.
- [9] I. Benjamini and N. Curien. Ergodic theory on stationary random graphs. *Electron. J. Probab.*, 17:no. 93, 20, 2012.
- [10] I. Benjamini and N. Curien. Recurrence of the \mathbb{Z}^d -valued infinite snake via unimodularity. *Electron. Commun. Probab.*, 17:no. 1, 10, 2012.
- [11] I. Benjamini, N. Curien, and A. Georgakopoulos. The Liouville and the intersection properties are equivalent for planar graphs. *Electron. Commun. Probab.*, 17:no. 42, 5, 2012.
- [12] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Critical percolation on any nonamenable group has no infinite clusters. *Ann. Probab.*, 27(3):1347–1356, 1999.
- [13] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Group-invariant percolation on graphs. *Geom. Funct. Anal.*, 9(1):29–66, 1999.
- [14] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Uniform spanning forests. *Ann. Probab.*, 29(1):1–65, 2001.
- [15] I. Benjamini, R. Lyons, and O. Schramm. Percolation perturbations in potential theory and random walks. In *Random walks and discrete potential theory (Cortona, 1997)*, Sympos. Math., XXXIX, pages 56–84. Cambridge Univ. Press, Cambridge, 1999.
- [16] I. Benjamini, R. Lyons, and O. Schramm. Unimodular random trees. arXiv:1207.1752, 2012.
- [17] I. Benjamini, E. Paquette, and J. Pfeffer. Anchored expansion, speed, and the hyperbolic Poisson Voronoi tessellation. arXiv:1409.4312.
- [18] I. Benjamini and O. Schramm. Harmonic functions on planar and almost planar graphs and manifolds, via circle packings. *Invent. Math.*, 126(3):565–587, 1996.
- [19] I. Benjamini and O. Schramm. Random walks and harmonic functions on infinite planar graphs using square tilings. *Ann. Probab.*, 24(3):1219–1238, 1996.
- [20] I. Benjamini and O. Schramm. Percolation in the hyperbolic plane, 2000.
- [21] I. Benjamini and O. Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6:no. 23, 1–13, 2001.
- [22] I. Benjamini and O. Schramm. Percolation beyond \mathbb{Z}^d , many questions and a few answers [mr1423907]. In *Selected works of Oded Schramm. Volume 1, 2*, Sel. Works Probab. Stat., pages 679–690. Springer, New York, 2011.
- [23] I. Biringer and J. Raimbault. The topology of invariant random surfaces. <http://arxiv.org/abs/1411.0561>.
- [24] L. Bowen. Periodicity and circle packings of the hyperbolic plane. *Geom. Dedicata*, 102:213–236, 2003.
- [25] L. Brouwer. On the structure of perfect sets of points. In *KNAW, Proceedings*, volume 12, pages 1909–1910.
- [26] R. M. Burton and M. Keane. Density and uniqueness in percolation. *Communications in mathematical physics*, 121(3):501–505, 1989.
- [27] N. Curien. A glimpse of the conformal structure of random planar maps. arXiv:1308.1807, 2013.
- [28] N. Curien. Planar stochastic hyperbolic infinite triangulations. arXiv:1401.3297, 2014.
- [29] N. Curien, L. Ménard, and G. Miermont. A view from infinity of the uniform infinite planar quadrangulation. *ALEA Lat. Am. J. Probab. Math. Stat.*, 10(1):45–88, 2013.
- [30] J. Ding, J. R. Lee, and Y. Peres. Markov type and threshold embeddings. *Geometric and Functional Analysis*, 23(4):1207–1229, 2013.
- [31] R. J. Duffin. The extremal length of a network. *J. Math. Anal. Appl.*, 5:200–215, 1962.
- [32] G. Elek. On the limit of large girth graph sequences. *Combinatorica*, 30(5):553–563, 2010.
- [33] G. Elek and G. Lippner. Sofic equivalence relations. *J. Funct. Anal.*, 258(5):1692–1708, 2010.

- [34] A. Gandolfi, M. Keane, and C. Newman. Uniqueness of the infinite component in a random graph with applications to percolation and spin glasses. *Probability Theory and Related Fields*, 92(4):511–527, 1992.
- [35] C. Garban. Quantum gravity and the KPZ formula [after Duplantier-Sheffield]. *Astérisque*, (352):Exp. No. 1052, ix, 315–354, 2013. Séminaire Bourbaki. Vol. 2011/2012. Exposés 1043–1058.
- [36] J. R. Gilbert, J. P. Hutchinson, and R. E. Tarjan. A separator theorem for graphs of bounded genus. *Journal of Algorithms*, 5(3):391–407, 1984.
- [37] J. T. Gill and S. Rohde. On the Riemann surface type of random planar maps. *Rev. Mat. Iberoam.*, 29(3):1071–1090, 2013.
- [38] O. Gurel-Gurevich and A. Nachmias. Recurrence of planar graph limits. *Ann. of Math. (2)*, 177(2):761–781, 2013.
- [39] O. Gurel-Gurevich and A. Nachmias. Recurrence of planar graph limits. *Ann. of Math. (2)*, 177(2):761–781, 2013.
- [40] O. Häggström. Infinite clusters in dependent automorphism invariant percolation on trees. *Ann. Probab.*, 25(3):1423–1436, 1997.
- [41] O. Häggström and Y. Peres. Monotonicity of uniqueness for percolation on Cayley graphs: all infinite clusters are born simultaneously. *Probab. Theory Related Fields*, 113(2):273–285, 1999.
- [42] A. Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- [43] Z.-X. He and O. Schramm. Fixed points, Koebe uniformization and circle packings. *Ann. of Math. (2)*, 137(2):369–406, 1993.
- [44] Z.-X. He and O. Schramm. Hyperbolic and parabolic packings. *Discrete Comput. Geom.*, 14(2):123–149, 1995.
- [45] T. Hutchcroft. Interacements and the wired uniform spanning forest. arXiv:1512.08509, 2015.
- [46] T. Hutchcroft. Wired cycle-breaking dynamics for uniform spanning forests. arXiv:1504.03928, 2015.
- [47] T. Hutchcroft and A. Nachmias. Uniform spanning forests of planar graphs. <http://arxiv.org/abs/1603.07320>.
- [48] V. Kaimanovich. Boundary and entropy of random walks in random environment. *Prob. Theory and Math. Stat*, 1:573–579, 1990.
- [49] V. A. Kaimanovich. Hausdorff dimension of the harmonic measure on trees. *Ergodic Theory and Dynamical Systems*, 18(03):631–660, 1998.
- [50] V. A. Kaimanovich. Random walks on sierpiński graphs: hyperbolicity and stochastic homogenization. In *Fractals in Graz 2001*, pages 145–183. Springer, 2003.
- [51] V. A. Kaimanovich, Y. Kifer, and B.-Z. Rubshtein. Boundaries and harmonic functions for random walks with random transition probabilities. *Journal of Theoretical Probability*, 17(3):605–646, 2004.
- [52] V. A. Kaimanovich and A. M. Vershik. Random walks on discrete groups: boundary and entropy. *The annals of probability*, pages 457–490, 1983.
- [53] V. A. Kaimanovich and W. Woess. Boundary and entropy of space homogeneous markov chains. *Annals of probability*, pages 323–363, 2002.
- [54] M. Krikun. Local structure of random quadrangulations. arXiv:math/0512304, 2005.
- [55] S. K. Lando and A. K. Zvonkin. *Graphs on surfaces and their applications*, volume 141 of *Encyclopaedia of Mathematical Sciences*. Springer-Verlag, Berlin, 2004. With an appendix by Don B. Zagier, Low-Dimensional Topology, II.
- [56] G. F. Lawler. A self-avoiding random walk. *Duke Math. J.*, 47(3):655–693, 1980.
- [57] R. J. Lipton and R. E. Tarjan. Applications of a planar separator theorem. *SIAM J. Comput.*, 9(3):615–627, 1980.
- [58] R. Lyons, B. J. Morris, and O. Schramm. Ends in uniform spanning forests. *Electron. J. Probab.*, 13:no. 58, 1702–1725, 2008.
- [59] R. Lyons and Y. Peres. *Probability on Trees and Networks*. Cambridge University Press, 2015. In preparation. Current version available at <http://mypage.iu.edu/~rdlyons/>.
- [60] R. Lyons, Y. Peres, and O. Schramm. Minimal spanning forests. *Ann. Probab.*, 34(5):1665–1692, 2006.
- [61] R. Lyons and O. Schramm. Indistinguishability of percolation clusters. *Ann. Probab.*, 27(4):1809–1836, 1999.
- [62] C. M. Newman and L. S. Schulman. Infinite clusters in percolation models. *J. Statist. Phys.*, 26(3):613–628, 1981.

- [63] R. Pemantle. Choosing a spanning tree for the integer lattice uniformly. *Ann. Probab.*, 19(4):1559–1574, 1991.
- [64] I. Richards. On the classification of noncompact surfaces. *Transactions of the American Mathematical Society*, 106(2):259–269, 1963.
- [65] K. Stephenson. *Introduction to circle packing*. Cambridge University Press, Cambridge, 2005. The theory of discrete analytic functions.
- [66] B. Virág. Anchored expansion and random walk. *Geom. Funct. Anal.*, 10(6):1588–1605, 2000.
- [67] D. B. Wilson. Generating random spanning trees more quickly than the cover time. In *Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996)*, pages 296–303. ACM, New York, 1996.

OMER ANGEL, TOM HUTCHCROFT
DEPARTMENT OF MATHEMATICS, UNIVERSITY OF BRITISH COLUMBIA
Email: {angel,thutch}@math.ubc.ca

ASAF NACHMIAS
SCHOOL OF MATHEMATICAL SCIENCES, TEL AVIV UNIVERSITY
Email: asafnach@post.tau.ac.il

GOURAB RAY
STATISTICAL LABORATORY, UNIVERSITY OF CAMBRIDGE
Email: rg508@cam.ac.uk